

12
TECHNICAL REPORT 81-1

THE IMPACT OF INFORMATION ON DECISIONS: COMMAND AND CONTROL SYSTEM EVALUATION


DECISION SCIENCE CONSORTIUM, INC.

Marvin S. Cohen
Anthony N.S. Freeling

February 1981

AD A 097 220

DTIC
ELECTE
S
ANN 2 1981


**Command & Control
Decision & Forecasting
Systems Program**

DEFENSE SCIENCES OFFICE • CYBERNETICS TECHNOLOGY DIVISION
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

INTRODUCTION
APPROVAL

81-4-2-152

12
LEVEL II

TECHNICAL REPORT 81-1

THE IMPACT OF INFORMATION ON DECISIONS:
COMMAND AND CONTROL SYSTEM EVALUATION

by

Marvin S. Cohen, and Anthony N. S. Freeling

Sponsored by

Defense Advanced Research Projects Agency
Under Contract MDA903-80-C-0194
DARPA Order No. 3831

Under Subcontract from

Decisions and Designs, Inc.

February 1981

THE VIEWS AND CONCLUSIONS CONTAINED IN THIS DOCUMENT ARE THOSE OF THE AUTHOR AND SHOULD NOT BE INTERPRETED AS NECESSARILY REPRESENTING THE OFFICIAL POLICIES, EITHER EXPRESSED OR IMPLIED, OF THE DEFENSE ADVANCED RESEARCH PROJECTS AGENCY OR THE UNITED STATES GOVERNMENT.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

DECISION SCIENCE CONSORTIUM, INC.
7700 Leesburg Pike, Suite 421
Falls Church, Virginia 22043
(703) 790-0510

DTIC
ELECTE
S APR 2 1981 D
D

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER TR-81-1	2. GOVT ACCESSION NO. AD-A097 220	3. RECIPIENT'S CATALOG NUMBER (9)
4. TITLE (and Subtitle) The Impact of Information on Decisions: Command and Control System Evaluation		5. TYPE OF REPORT & PERIOD COVERED Technical report
6. AUTHOR(s) Marvin S. Cohen Anthony N. S. Freeling		7. PERFORMING ORG. REPORT NUMBER Technical Report 81-1
8. PERFORMING ORGANIZATION NAME AND ADDRESS Decision Science Consortium, Inc. 7700 Leesburg Pike, Suite 421 Falls Church, Virginia 22043		9. CONTRACT OR GRANT NUMBER(s) MDA903-80-C-0194, DARPA 79-312-0729 (Sub) Order-22
10. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209		11. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Darpa Order No. 3831
12. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		13. REPORT DATE February 81
		14. NUMBER OF PAGES 99 (2) 128
		15. SECURITY CLASS. (of this report) Unclassified
		16a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for publication release: distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Decision analysis Management Information Systems Value of information Data base system design Command, Control, Communications, and Intelligence Indications and Warning		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) See reverse side.		

In work reported here, Decision Science Consortium, Inc. (DSC) has examined the application of the decision analytic concept of value of information to the design of information systems. A common characteristic of systems in current use is that they often provide vast quantities of partially relevant data, while failing to identify the information which the decision maker actually needs to solve his problem.

Current evaluation techniques for information systems appear to bypass this problem altogether. The concept of Value of Information (VOI), however, implies that information has value to the extent that it can alter decisions and improve payoffs. Yet, the application of VOI techniques, as they now stand, to information systems is prohibitively complex and rests on implausible assumptions. The aim of the present work is to devise modifications of standard VOI techniques which make them simple enough and realistic enough to apply to information system design, while retaining a basic reference to the impact of information on decisions.

- (1) Explicit reference to all the information in a data base can be omitted from a VOI analysis if acts are modeled as events (Brown, 1975). Assessment heuristics are presented which facilitate modeling acts as events in this (or any other) context.
- (2) Non-optimal uses of information, and the evaluation of decision aids, are also accommodated by means of these heuristics.
- (3) The omission of critical uncertainties from a VOI analysis is examined within the acts as events framework.
- (4) A method for omitting specified options from a VOI model is presented.
- (5) A multiattribute decomposition into information value and usability is derived.

Finally, concepts are outlined for the programming of "intelligent" real-time information selection for a particular user.

Unclassified

SUMMARY

In work reported here, Decision Science Consortium, Inc. (DSC) has examined the application of the decision analytic concept of value of information to the design of information systems. Automated data base systems play an increasingly prominent role in a variety of areas - including Command, Control, Communications, and Intelligence (C³I), Indications and Warning (I & W), and business management. However, a basic problem of data base design has not been solved: what information should be included in the system, and what subset of that information should be presented to a user, so as to best achieve the objectives of the relevant organization? A common characteristic of systems in current use is that they often provide vast quantities of partially relevant data, while failing to identify the information which the decision maker actually needs to solve his problem.

Current evaluation techniques for information systems appear to bypass this problem altogether. Evaluation in terms of data-processing parameters, like channel capacity or memory size, ignores the ultimate objectives of the system and seems to assume, simply, that more information is better. Direct assessment of information quality, in terms of such attributes as relevance and accuracy, fails to ensure that the actual impact of information on decisions (hence, on ultimate objectives) is considered. Multi-attribute utility models similarly have not explicitly required consideration of how information is used in decision making.

The concept of Value of Information (VOI) implies that information has value to the extent that it can alter decision and improve payoffs. However, the application of VOI techniques, as they now stand, to information systems is prohibitively complex. These techniques presuppose a highly structured decision problem, in which information, the uncertainties to which it pertains, and

the options available to the decision maker are all specified in detail. Moreover, they assume that the decision maker will behave optimally in the light of the information he receives. Complex, multipurpose information systems, on the other hand, are expected to operate in a variety of environments, some of which cannot be predicted in advance. And, they must serve users who cannot always conform to established normative ideals.

The aim of the present work is to devise modifications of standard VOI techniques which make them simple enough and realistic enough to apply to information system design, while retaining a basic reference to the impact of information on decisions. In doing so, we proceed by steps of (roughly) increasing simplification:

- (1) Explicit reference to all the information in a data base can be omitted from a VOI analysis if acts are modeled as events (Brown, 1975). Assessment heuristics are presented which facilitate modeling acts as events in this (or any other) context. By means of these heuristics, a system designer can examine tradeoffs between the information value of a system and its usability.
- (2) The same heuristics allow the designer to accommodate non-optimal uses of information and to evaluate decision aids which support inferential and decision making processes.
- (3) The manner in which information affects judgments concerning critical events need not be explicitly modeled. It is shown that modeling acts as events may reduce the cost, in terms of credibility, of omitting this part of a VOI analysis. A convenient form of assessment, when critical events are not modeled, is in terms of the expected cost of errors (or opportunity loss). The application of this notion when acts are modeled as events, however, raises special difficulties, which are dealt with.

(4) If certain assumptions are acceptable, the options facing a decision maker need not be specified in a VOI analysis. Information may be evaluated in terms of the overall probability that it will cause a decision maker to switch from an otherwise preferred option, and the expected swing in utility if he does so.

(5) The information value and usability of a system can be distinguished, within this evaluation technique, by decomposing the probability of switching options. The result is a multiattribute utility model which has well-defined attributes and a well-motivated rule for combining them, and which refers explicitly to the impact of information on decisions. This approach extends, once again, to the modeling of non-optimal behavior and to the evaluation of inference and decision aids.

We turn next to the application of these concepts to the interaction with a system by a particular user. The objective is to program into the data base "intelligent" real-time information selection for the user. We outline an interactive procedure which focuses the user's attention on the potential decisional impact of information and which guides him to the data categories of most pertinence. Again, a variety of levels of simplification are explored. Procedures for mapping critical events, about which the user is uncertain, onto data categories are also examined.

TABLE OF CONTENTS

	Page
SUMMARY	ii
1.0 INTRODUCTION	1-1
1.1 The Need for Goal-oriented Information System Design	1-1
1.2 The Information System Evaluation Problem . .	1-2
1.3 The Concept of Value of Information and Current Techniques	1-3
1.4 Preliminary Distinctions	1-5
1.5 Outline of Report	1-10
2.0 BACKGROUND RESEARCH	2-1
2.1 Non-decision Oriented Approaches	2-1
2.1.1 Data-processing parameters	2-2
2.1.2 Information quality	2-3
2.1.3 Multiattribute utility	2-7
2.1.4 User information selection	2-9
2.2 Value of Information	2-10
2.2.1 A decision problem	2-10
2.2.2 Value of perfect information	2-12
2.2.3 Opportunity loss	2-12
2.2.4 Value of imperfect information	2-13
2.2.5 Non-negativity of information value. .	2-16
2.2.6 Non-decisional impact of information .	2-17
2.2.7 Conditions of positive value	2-19
2.2.8 Application of VOI to information systems: practicality	2-20
2.2.9 The problem of specifying structure. .	2-22
2.2.10 Standard simplifications of VOI . . .	2-22
2.2.10.1 Opportunity loss	2-23
2.2.10.2 Perfect information	2-23
2.2.10.3 Scenario sampling	2-24

	Page
2.3 Modifications in VOI Techniques	2-26
2.3.1 The consistency condition	2-26
2.3.2 Acts as events	2-27
2.3.2.1 Failures of consistency	2-27
2.3.2.2 Sufficient modeling	2-30
2.3.3 Value of analysis	2-31
2.3.3.1 Watson and Brown	2-32
2.3.3.2 Tani	2-36
2.3.3.3 Nickerson and Boyd	2-37
3.0 TECHNIQUES FOR INFORMATION SYSTEM DESIGN	3-1
3.1 Unspecified Experimental Outcomes	3-2
3.2 Acts as Events in System Evaluation	3-3
3.3 Assessment heuristics	3-10
3.3.1 Auxiliary decision tree	3-10
3.3.2 Assessment of act probabilities	3-12
3.3.3 Assessing c	3-14
3.3.4 Assessment of expected utilities	3-15
3.4 Non-optimal Behavior	3-17
3.4.1 Application to system design	3-19
3.4.2 Application to value of analysis	3-20
3.5 An Inferential Structure for VOI	3-21
3.5.1 Inference tree representation	3-22
3.5.2 Computation of VOI	3-22
3.5.3 Assessments on u; independence assumptions	3-24
3.6 Unspecified States of the World: Credibility	3-27
3.7 Utility Swing	3-30
3.8 Unspecified Options	3-35
3.9 Decomposing Shift Probabilities: Information Value and Usability	3-36
3.10 Application to VOA	3-39

	Page
4.0 INFORMATION EVALUATION FOR SYSTEM USERS	4-1
4.1 "Intelligent" Computer Aids for Information Selection	4-1
4.2 Outline of Proposed Program	4-2
4.3 Levels of Analysis	4-4
4.4 Presentation of Data	4-6
5.0 TECHNICAL NOTE: CREDIBILITY	
6.0 REFERENCES	
APPENDICES	
DISTRIBUTION LIST	

1.0 INTRODUCTION

1.1 The Need for Goal-oriented Information System Design

There is a curious paradox in the current status of automated information systems - whether in Command, Control, Communications, and Intelligence (C³I), Indications and Warning (I & W), or business management. On the one hand, as Zani (1979) notes regarding management information systems (MIS), they "have not really been designed at all. They have been spun off as by-products of automating or improving existing systems...." It might be supposed that such a bottom-up approach would result in a fairly painless infiltration of information systems into pre-existing organizations. Yet, this has been far from the case. In many instances, the introduction of systems in this way has produced controversy and failed to satisfy initial expectations (e.g., Miller, 1980; Beard, 1977; Swanson, 1974).

A consensus seems to be emerging that a more active design process is required if information systems are to be optimally exploited:

- Andriole (1980) argues that the focus in C³ on "timely, rapid, survivable, and secure" delivery of information reflects lack of appreciation for what is done with the information after it is delivered, i.e., how commanders cognitively process it in order to make decisions.
- Gorry and Morton (1971) describe how management information systems have focused on techniques for automating routine operations rather than support of problem-solving and strategic planning.
- Mintzberg (19__) argues that management information systems tend either to flood users with data or else summarize and average to the point of blandness. There is no intelligent selection of the details needed for decision-making.

- Shlaim (1976) stresses that analysis and interpretation have lagged behind the sheer production of data in the area of monitoring and warning.
- Miller (1980) cites severe organizational and institutional resistance to the implementation of command and control projects.

As a result of such considerations, the Deputy Assistant Secretary of Defense has recently urged that Command and Control systems be designed and evaluated explicitly in terms of their contribution to mission success (Von Trees, 1980). Zani (1979) similarly concludes that MIS design should begin with a fundamental analysis of management decision functions. Finally, Daly and Andriole (1979) stress that Indications and Warning system design must refer to the purposes for which warnings are sought.

1.2 The Information System Evaluation Problem

Unfortunately, the implementation of a goal-oriented design process for information systems runs into technical problems. Systems for the management and display of information should help decision makers achieve their objectives by making better decisions. In contrast to transportation systems or weapon systems, information systems do not (necessarily) change the world in which action occurs, but have their principle impact on the cognitive processes that lead to action. Yet a conspicuous gap in the technology of design is an understanding of the relationship between information and decisions. As a result, evaluation of information systems has typically omitted reference to the decisions the systems are presumably designed to support.

In general, a viable technique for assessing systems must be both practical and relevant. It must involve attributes which are realistically measurable and which at the same time reflect

a system's potential contribution to ultimate objectives. In the case of information systems, practicality is frequently achieved by measures which ignore this distinctive character of information, and whose relevance is therefore doubtful.

1.3 The Concept of Value of Information and Current Techniques

In the work reported here we take the opposite tack. We start with an analysis which focuses on the decisional impact of information, although its shortcomings from the point of view of practicability are rather severe. We then explore ways in which the analysis can be simplified and modified in order to arrive at an operative evaluation technique for the data base content of information systems.

The starting point suggested by this strategy is the decision analytic concept of Value of Information (VOI). The essence of the VOI concept is that information has value only to the extent that it can alter decisions, with a resultant change in expected payoffs. As we have implied, the direct application of current VOI techniques presents problems, since even the evaluation of a single data item in a single scenario may require an unmanageable number of assessments (Raiffa and Schlaifer, 1961). Quite apart from practicality, however, there are more troubling and fundamental objections to the straightforward use of VOI as it now stands.

VOI applies to the value of performing an experiment. The experiment is expected to reduce uncertainty concerning an event relevant to the decision maker's choice of action. In order to assess VOI, therefore, one must have specified the experiment, the uncertainty to be reduced, the potential bearing of the experiment on the uncertainty, and the choices which may be affected. In other words, the problem must be highly structured. A second assumption is that the decision maker's choices subsequent to

receipt of the information will be rationally predictable on the basis of that structure.

These demands are not typically satisfied:

- Information systems will be expected to operate in a very large and not very well specified set of scenarios. C³ systems, for example, must deal with contingencies determined in part by our adversaries. Options and major uncertainties may not be known in advance.
- Users may diverge from established normative theories in the manner in which they handle information. Political factors quite unrelated to organizational goals often distort the flow of information in organizations (Huber, 1980a, 1980b). Limitations of time, cognitive capacity, and knowledge prevent individuals from making the normatively correct use of information presupposed by VOI.

A corollary of the fact that users do not always behave optimally is that information systems may help them to behave more optimally. Such systems may in fact provide assistance in structuring the problem, i.e., in generating options and identifying important parameters; as well as in drawing inferences and prioritizing actions. Traditional VOI techniques do not apply directly to the evaluation of this type of information.

A common element in these considerations has been the question of appropriate structure. The structure presupposed by traditional VOI may be unknown in advance, disregarded by the decision maker at the time of action, or supplied to him at that time by the information system. This common element is our starting point. Generalizations of VOI which handle it might produce at the same time a significant simplification of the required assessments

and computations. The objective of the research reported here is to pursue that possibility, in order to discover practicable evaluation techniques for information systems which nonetheless retain the VOI concept with its explicit reference to decisions.

1.4 Preliminary Distinctions

It will be useful at this point to set boundaries on our present concerns and to mark some distinctions which will figure in the application of VOI to information systems.

(1) First, the items to be evaluated are the content of the information system (and to a lesser degree the manner in which information items are displayed), in contrast to the hardware and software configuration.

(2) A truly general evaluation method, however, will allow these items to be either experiments or facts. An experiment is an observation which has a number of possible outcomes. For example, the application of VOI to a forecasting model involves the issue of what indicators to include, where each indicator is a variable which can assume a range of levels. Experiments play a role in other contexts as well--e.g., in command and control, the value of different kinds of information about the location and identity of hostile platforms in the immediate vicinity.

Facts, on the other hand, are constants and are (in principle) knowable in advance of their selection for use in a decision context. The particular outcome of a previously performed experiment is a fact. Facts may concern past events (e.g., international crises), but also include theoretical and historical generalizations.

Some data base systems consist exclusively of facts; an example is Executive Aids for Crisis Management (Mahoney, et al., 1978; Spector, et al., 1978) which consists of historical facts about international crises organized according to actions, objectives, problems, and other descriptors. More usually, there is a mix of facts and experiments. The designer of a combat center, may decide to include within the data base continuously updated information on threat locations, as well as standing intelligence on threat capabilities.

Many facts to be considered for inclusion in an information system will be initially unknown to the system designers (e.g., the specifics of threat capabilities). The decision of whether or not to include these facts in the data base must, therefore, involve two phases: first, whether to expend the resources required to make them available; second, whether or not to include the output of the first phase (known facts) in the data base system. "Unknown facts" may, in the initial phase, either be characterized in terms of a variable the actual level of which is uncertain (e.g., weapon range)--or else by means of a more generic category (e.g., "the events leading up to the Suez Crisis"). In the former case, unknown facts are treated analogously to experiments.

In general, experiments or facts may be considered one by one for exclusion or inclusion in a data base--or else they may be grouped into larger categories. At the limit, the entire contents of a proposed data base can be evaluated as a single unit. Grouping of items will be necessary when the items interact in their impact on decisions--as, for example, when one item cannot be interpreted properly in the absence of another.

(3) Another important distinction among information items concerns the level of analysis which they represent. Certain

items may provide higher-order information regarding the implications of other data for judgment and/or action. For example, in EWAMS (IPPRC, 1979; DCI, 1978) the number of bilateral interactions of a given type between two countries serves as the basis for an inference regarding the likelihood of a crisis. In this case, both levels of analysis are available to the user: datum and inference.

As another example, in submarine command and control raw bearing measurements are used to estimate the range of a target. Target range may in turn figure (together with facts about adversary capabilities) in still higher-order inferences concerning the probability of being within threat weapon range (Cohen and Brown, 1980). A still higher-order inference might regard the action (e.g., fire now or continue to approach) which maximizes expected utility. An important issue in the design of such a system is the selection of levels to be presented: To what degree does the user require evidence as well as conclusions? To what degree will the presentation of conclusions mark an improvement over the inferences the user would have drawn on his own from the evidence?

(4) A different kind of distinction concerns the time at which information is evaluated. We consider that value of information can be computed and information choices made at either (or both) of two stages: system design or system use. We begin by assuming a hypothetical "Universal Data Base" containing all information items which are of any relevance whatsoever. (This is a rather loosely defined and open-ended set, since it contains all possible inferences from its members.) The first stage of selection (Figure 1-1) requires the selection of a subset of the Universal Data Base to serve as the "Actual Data Base" for a given system. This selection is performed by the designers of the system with its intended users in mind. The second stage of selection produces a subset of the Actual Data Base to serve as a "Virtual




Figure 1-1
Levels of Information Selection

Data Base" tailored to the requirements of a current decision in a specific scenario (Figure 1-1). This stage directly involves a user of the system (although provision for such specification is a task for the system designer).

Information evaluation techniques may be applied at either (or both) of these stages. We refer to the interaction of the designer with such a technique as the "designer dialogue", and the interaction of the user as the "user dialogue". Somewhat different constraints affect choices at these two levels. When the user of a C² system selects a virtual data base (e.g., the contents of a display screen or of core memory), his major limitations are likely to be time and cognitive capacity. Other considerations might include the number of display surfaces and the personnel available to record and analyze the information selected.

The system designer must consider these same factors in a general way, especially if a user-controlled stage of selection is not provided for. But additional factors enter into the designer dialogue, e.g., limits on the size of long-term storage devices, and the costs for research and development.

(5) A final demarcation of our interest concerns how information items are to be evaluated. Information has many effects: it may simply satisfy curiosity, and it may increase confidence in one's previous decisions or in those of someone else. These effects are certainly not to be ignored. But our central focus here, as implied by our interest in VOI, will be on information's decisional impact. The measures which we develop can, we assume, be combined subsequently with other aspects of value, within the context of a general multi-attribute utility model.

1.5 Outline of Report

Chapter Two examines information system evaluation techniques in current practice. It consists of three parts. In the first, we look at techniques which seem practicable but which do not explicitly refer to decision making. In the second, we examine the VOI concept and shortcomings of current VOI techniques. And, in the third, we describe some recent modifications of VOI which hold promise for the application to information systems. These modifications were originally designed to handle non-optimal or not fully modeled decision making (Brown, 1975) and the evaluation of decision analysis (Brown and Watson, 1975).

Chapter Three reports the major results of the study. It turns out that the modifications of VOI mentioned above make a significant contribution to the problem of applying VOI to unspecified scenarios and, in general, to the goal of a simple, practical evaluation technique in system design. Some additional modifications are suggested to further simplify the technique and reduce the degree of specification required. These modifications in turn are shown to contribute to the methodology for evaluating decision analysis and for representing non-optimal behavior.

Chapter Four applies the foregoing results to the on-line selection of information by system users. A considerable increase in system flexibility and in accommodation to individual differences can be achieved by applying VOI concepts to the interaction between user and system in a particular decision problem, rather than merely at the system design stage.

Finally, an Appendix explores some possible applications of fuzzy set theory to the problem of information evaluation.

2.0 BACKGROUND RESEARCH

Two purposes motivate the development of an evaluation methodology applicable to information systems. Such a methodology can help set priorities for the allocation of resources among diverse options - e.g., command and control, weapons, and force levels. On the other hand, it can lead to the improvement of information system design for a given expenditure of resources. In both cases, the role of an evaluation technique is in the measurement of a proposed system's contribution to ultimate objectives, or utility.

Currently practiced techniques differ in the features of information systems which serve as indicants of ultimate value. As a consequence, they differ (a) in the clarity of the relationship between such features and utility, and (b) in the number and difficulty of the required assessments. A brief survey of available techniques will show that the objectives of validity and practicability tend to conflict, and that it is difficult to achieve an acceptable level on both at once.

2.1 Non-Decision Oriented Approaches

In this section we will briefly survey four approaches to information system evaluation. None of these approaches involves explicit reference to decisions. They differ according to the avowed criteria of evaluation:

- Data-processing parameters
- Information quality
- Multi-attribute utility
- User information selection

2.1.1 Data-processing parameters. In the area of C^3I , communications technology has traditionally taken the lead. Evaluation of C^3I systems has tended to focus on properties of information transmission and storage. The availability of well-developed theories in this area has encouraged the use of measures like channel capacity, connectivity, memory size, and computation speed. In some methodologies (e.g., TRI-TAC, described in Miller, 1980), these measures are combined by a weighting scheme into a single index of communication performance.

Unfortunately, the measure of worth deriving from this approach makes no reference to the real objective of the system: viz., mission accomplishment. To assume that communication performance reflects such an objective in a straightforward, or even monotonic, fashion is to accept on faith that "more information is better". This is to ignore human cognitive and organizational constraints which dictate a need for information filtering and decision aiding (Andriole, 1980).

Moreover, it is hard to see how meaningful weights could be assigned, and an appropriate balance struck, between such competing claims as memory size and computation speed without reference to the uses of the system. For example, a system designed to calculate the trajectory of an approaching weapon would require a different tradeoff on these dimensions than one designed to evaluate the likelihood of a political coup in Iraq. "Communication performance" itself cannot be assessed in abstraction from a mission.

Measures of this sort may prove useful (if appropriately linked to ultimate objectives) for the evaluation of alternative hardware or software designs. But they do not apply at all to the selection of the content of the information to be provided. For this reason, it does not seem desirable to construct an

evaluation technique with these measures as a basis. It has been suggested (Alberts, 1980) that an evaluation technique might start with data-processing parameters, but employ "linkage models" which describe the relationship between these parameters and higher-level indicants of system utility. Such a technique, while acknowledging the importance of ultimate utility, would have to ignore some of the most critical sources of variation among proposed systems--i.e., the nature of the information presented.

The optimal exploitation of more effective weapons and more highly mobile forces requires C³I systems capable of supporting rapid, accurate decision making. The advantage in an engagement may belong to the side which can saturate the command and control capacity of an adversary. Yet current C³I systems tend to present large quantities of data about the environment, adversaries, own unit, and weapons, in relatively raw undigested form (cf., Cohen and Brown, 1980). An exclusive stress on communications properties, by deflecting attention from what is done with the information communicated, is unlikely to lead to dramatic improvement.

2.1.2 Information quality. A second approach evaluates information systems at a higher level of abstraction and has been particularly prominent in the literature on Management Information Systems (MIS). Information content and information delivery are characterized in terms of such dimensions as accuracy, relevance, timeliness, clarity, and readability (Swanson, 1974; Gallagher, 1974). Direct subjective assessments of these properties may be obtained in a laboratory context or in large-scale prototype testing. Again, measures on these dimensions (or others defined in terms of them) may be combined by a weighting scheme into a single index of worth.

On the face of it, these properties seem to bear a close relation to the contribution of a system to ultimate objectives. The notion of precision, for example, is defined as the proportion

of the supplied information items which are deemed relevant by the user. Thus, weight is given to the need for an intelligent selection of the information to be provided. An adequate system must achieve a suitable balance between precision and completeness, defined as the proportion of relevant items which are supplied (cf., Cleverdon, 1962; Smith, 1972).

Similarly, requirements of usability are acknowledged in such dimensions as clarity and readability. Again, successful systems must strike a balance between usability and informativeness in the more abstract sense, represented by relevance, accuracy, and timeliness. A distinction of this sort between two classes of properties has become widely accepted (Herner and Snapper, 1978; Larcker and Lessig, 1980). For example, Smith (1972) proposed two sets of criteria, one dealing with efficiency from the operator viewpoint and one with effectiveness from the user viewpoint. In his efficiency set, Smith listed a variety of attributes, reflecting both cognitive and organizational factors: orientation toward a single organization (or, presumably, body of users), functional and technical integration with the target organization, uniqueness (or lack of overlap with other systems), flexibility, efficacy of processing procedures and programs, and efficacy of the man/machine interface. In his set of effectiveness attributes, Smith includes the following: relevance, accuracy, timeliness, sufficiency, conciseness, consistency of the data source, user confidence, and news or discovery value of messages.

A closer examination of this approach reveals, however, that the relation of criteria to utility has not been sufficiently clarified. Since the goal is to define the "perceived" value of information, justification of attributes which were initially posited a priori has been sought in studies of subjective judgments. For example, factor analyses of experimentally elicited

responses suggests that evaluations can be roughly organized in accordance with the proposed criteria (e.g., Zmud, 1978; Larcker and Lessig, 1980). Unfortunately, however, there is reason to suppose that subjective judgments tend to be guided by considerations other than ultimate utility.

Often a respondent is simply invited to indicate which data elements or information are "relevant" to him. In this case, no direction is given as to what the information should be relevant for. Data which have no effect on decision making (because, for example, they are already known to the user) might be counted "relevant" if they are related to the topic of concern. In some studies the respondent may be asked to check only those elements which are relevant in that they are likely to be "used". On occasion, the instructions request the respondent to check off items that are to be "used for decision making." The latter approach is in fact seldom employed. But even when it is used, critical ambiguities as to what is being assessed remain unresolved. Information may be used in decision making and yet have a very low probability of changing the decision and affecting utility. Moreover, no consideration is given to the relative importance of the decisions which might be affected.

Little attention has been paid within this approach to rules for weighting and combining attributes into a single index of utility (see Keeney and Raiffa, 1976). Factor analytic methods for extracting subjective dimensions from evaluative judgments (Zmud, 1978; Larcker and Lessig, 1980) implicitly impute orthogonality to the extracted dimensions. Yet in terms of utility, additive relationships do not always seem plausible. For example, a decrease in usability cannot always be fully compensated by an increase in informativeness (e.g., accuracy). If usability falls to zero (e.g., in the case of an illegible display), utility is presumably zero regardless of accuracy. This suggests a multiplicative relationship. Thus, while greater

accuracy may be achievable only by sacrificing usability, in terms of utility the more of one, the more valuable is a given level of the other. A similar argument applies to timeliness and accuracy. More time may buy more accuracy. But a system that never produces information on time is worthless, regardless of its accuracy. If orthogonality is to be imputed to these dimensions on the basis of factor analytic results, we can only conclude (once again) that respondents were not assessing utility.

The problems with "relevance" noted above spill over, of course, to derived measures like precision and completeness, which are defined in terms of relevance. But these measures raise fundamental questions about system evaluation in their own right. Completeness, as Cooper (1976) notes, places heavy stress on the presumed cost of "unexamined documents", i.e., relevant information which is not supplied by the system. But it is far from clear that this is a well-defined, quantifiable set except in very special circumstances. To be sure, in a detection problem (e.g., signalling the presence of enemy platforms in an area) misses are quantifiable in principle and are quite relevant to system evaluation. But in general it is not clear what the set of "all facts" on a given topic might be, nor why this set should be considered in evaluating the utility of the facts that are supplied. As Cooper notes, information never received by the user cannot affect his decisions. (It may, but need not, affect our assessments of likelihoods and utilities for the outcomes of his decisions.)

In short, it is neither feasible nor justified to automatically penalize a system when it does not present a "relevant" item of information. The more direct and valid approach is to compare, in terms of utility, systems which supply a particular type or item of information with those which do not.

An important result of treating information selection as if it were a detection problem is to slight the value of inference and decision aids. Such aids are not easily reconciled with the paradigm of a fixed set of facts from which data is to be selected. On the contrary, they may assist the user in providing a structure for the problem within which such facts can be interpreted. In Management Information Systems, as in C³I, the emphasis so far has been on information systems which perform routine repetitive tasks, at the level of "operational control," rather than systems which can support problem-definition and problem-solving at the level of "strategic planning" (Gorry and Morton, 1971). Major advances in the latter area will come not from more "complete" coverage of data, but from enhanced information processing ability.

2.1.3 Multi-attribute utility. A considerable body of techniques has been developed (e.g., Keeney and Raiffa, 1976) for measuring utility when competing alternatives vary on numerous relevant dimensions. Multi-attribute utility (MAU) models specify how measurable properties of a system are related to overall system utility by means of mediating attributes and conditioning variables. Development of such a model consists of several stages. Overall utility is decomposed into subfactors, and these are further decomposed down to the level of measurable properties. Functions are assessed on the measurable properties to express their contribution to factors at the next higher level. Combination rules are employed to show how the value of a factor is determined by the subfactors subordinate to it, and importance weights are assigned to the subfactors.

MAU models provide a single index of utility for complex systems. They have been applied to such C³I components as the single channel ground and airborne radio system (Chinnis, et al., 1975). "Value diagrams," which may be regarded as variants of MAU

models, have been applied to the evaluation of alternative intelligence collection platforms in Barclay, Brown, et al., 1977. In this application, the value of a particular platform is expressed in terms of its contribution within different collection modes (photographic, radar, etc.), for different types of target, for different types of threat, in different geographical regions.

MAU models, in contrast to the data processing and information quality approaches, give explicit attention to utility and to the form in which attributes combine. The objective is to facilitate evaluation by breaking it down into simpler components. Nonetheless, there has been little or no attention to the particular nature of information systems. The assessment of information value is not decomposed in such a way as to make explicit reference to its impact on decisions. The validity of the assessments obtained thus depends on the assumption that the assessor takes this impact implicitly into account.

For example, in the evaluation of intelligence platforms, importance weights must be assigned to different collection modes for a given target type, threat type, and geographic region. There is no guarantee that in assigning these weights, proper regard is paid to the decisions that would be affected. Yet information has value in this context chiefly to the extent that it can alter decisions.

In sum, although the MAU methodology is suitable for application to complex systems in general, there is as yet no adequate specification of the criteria or attributes which are particularly appropriate for information systems. In light of the previously noted problems with such measures as relevance, completeness, and precision, specifying such attributes promises to be a non-trivial task.

2.1.4 User information selection. A final approach is to rely not on assessments, but on behavior. Information selection decisions may be observed as they are actually made. Analysis of such decisions, which are assumed to reflect user needs, can guide system design (cf., Davis, 1974). In the adaptive information selection (AIS) method developed by Perceptronics (Samet, et al., 1976, 1977), a model of the information choices by individuals or organizations is used to determine which items will be routed where. The model is capable of adapting to changing preferences for information as circumstances alter.

This approach is ideally suited for the design of systems intended to automate and replace existing procedures. AIS is likely to channel information in a way which is compatible with the capacities and preferences of its users. Nonetheless, for this very reason, it has important limitations. Perfect mirroring of current procedures is incompatible with the achievement of certain kinds of improvement. AIS may in fact perform better than the users which it models when their errors can be accounted for as random noise. But it cannot correct systematic biases, fallacies in reasoning, or mistaken assumptions. Moreover, it cannot lead to the provision of information - e.g., decision aids - which significantly enlarge present capabilities.

Behavioral data have a status somewhat comparable to holistic subjective assessments of information value. Intuitions of users concerning their needs are a valuable source of insight for the system designer. However, these data need to be supplemented by a more analytical consideration of what the data are needed for, i.e., what their impact is expected to be on subsequent decisions.

2.2 Value of Information


The standard decision analytic concept of the value of information is well described in most of the basic textbooks (e.g., Brown, Kahr, and Peterson, 1974; Raiffa, 1968), and is pursued in greater depth and complexity in more advanced work (e.g., Raiffa and Schlaifer, 1961). Although information can have intrinsic value by increasing knowledge, the thrust of VOI is to evaluate information gathering acts (or "experiments") in terms of their instrumental impact on subsequent decisions. (Our notation will be a simplified version of the notation used in Raiffa and Schlaifer.)

2.2.1 A decision problem. If he is unable to obtain further information, a decision maker might face a simple decision problem like the one depicted in Figure 2-1. In this problem, he must choose between two options, a_1 or a_2 . One of two events or states of the world, s_1 or s_2 , will turn out to be the case. Depending on which combination of act and event obtains, he will experience a utility, $u(a,s)$. In order to decide between a_1 and a_2 , the decision maker assesses utilities for the terminal node of every path through the tree. He also assesses probabilities for each event, $P(s_1)$ and $P(s_2)$. He can now choose between a_1 and a_2 by "averaging out and folding back" the tree. First, he computes the expected value of each action by taking the sum of the utilities for that action weighted by the probabilities:

$$u^*(a) = E_s u(a, \tilde{s}) = \sum_s P(s) u(a, \tilde{s}).$$

The expected value of the decision problem itself is the expected value of the best alternative:

$$u^* = \max_a E_s u(a, \tilde{s}).$$



A Simple Decision Problem
Figure 2-1

2.2.2 Value of perfect information. It is quite simple to assess the expected value of perfect information (EVPI) about $S = \{s_1, s_2\}$ for this decision maker. In effect, we "flip the tree," placing the uncertainty node for s before, rather than after, the decision node. We thus assume the decision maker has information about s when he makes the decision. Now, rather than selecting the option with the largest utility averaged across possible states of the world, he can select the option with the largest utility within each state of the world, whatever it turns out to be:

$$u^*(PI) = E_s \max_a u(a, \tilde{s}).$$

The expected value of perfect information is simply the value of the decision with the information less the value of the decision without it. Clearly,

$$E_s \max_a u(a, \tilde{s}) \geq \max_a E_s u(a, \tilde{s});$$

hence,

$$u^*(PI) \geq u^*.$$

Thus, the expected value of perfect information is positive or zero. (Note that this does not imply that the decision maker, even with perfect information about s , selects the option with the best outcome. s may not be the only relevant uncertainty, in which case $u(a, s)$ is not a realized utility but an expectation over the unmodeled events. Alternatively, $u(a, s)$ may be the expected value of a probability distribution assessed on utility u , for each a and s .)

2.2.3 Opportunity loss. Some insight into the dependence of the value of perfect information on decisions is obtained through an alternative formulation, in terms of expected opportunity loss (EOL). Let us assume that without any further information the decision maker prefers option a_1 (i.e., a_1 maximizes $E_s u(a, \tilde{s})$).

Now we can compare, for each s , the utility he receives from a_1 with the utility he would have received if he acted on perfect information about s . This difference is referred to as an "opportunity loss:"

$$l(a_1, s) = \max_a u(a, s) - u(a_1, s).$$

It can be shown that the expected value of perfect information is the same as the expected opportunity loss for a_1 :

$$\begin{aligned} \lambda^* &= E_s l(a_1, \tilde{s}) \\ &= E_s [\max_a u(a, \tilde{s}) - u(a_1, s)] \\ &= \text{EVPI}. \end{aligned}$$


Observe that for values of s in which the already preferred act, a_1 , is in fact the best choice, $\max_a u(a, s) = u(a_1, s)$, and there is no gain in utility from perfect information. Information has value only if it can change decisions. In effect, the value of perfect information is the expected cost of errors.

2.2.4 Value of imperfect information. Often a decision maker is able to obtain information about possible states of the world which, while not perfect, nonetheless has a bearing on his choice of action. Let us assume that before choosing between a_1 and a_2 , he has the chance to perform an experiment, e_1 , with possible outcomes z_1 and z_2 . The expected value of this experiment can be assessed by adding an uncertainty node, representing its possible outcomes, to the decision tree prior to the decision node. In Figure 2-2, we see that the decision maker can either choose immediately between a_1 and a_2 (after the dummy experiment, e_0) or else perform e_1 first.

A convenient assumption in analyzing value of information is that utility is a linear function of some measure (e.g., money) of total consequences (Raiffa and Schlaiffer, 1961; Lavalley, 1968).

Figure 2-2

Decision Tree for Value of Information



Then the utility assessed for each path through the tree can be expressed as two additive segments:

$$u(e, z, a, s) = u(a, s) - c(e),$$

where $c(e)$ represents the cost of performing the experiment. (We further assume that $c(e)$ is independent of experimental outcomes, z .) As a result, we can compute the value of the experiment e_1 (disregarding its cost) and then compare this value with $c(e_1)$, in order to decide whether to purchase it.

To evaluate the experiment e_1 , the decision maker must make some additional assessments: (a) the probability of receiving a particular observation from the experiment, $P(z|e_1)$ and (b) the probability of states of the world conditional upon these observations, $P(s|z, e_1)$. Usually, the most natural way to obtain these probabilities is to assess $P(s|e_1)$ and $P(z|s, e_1)$, and then to use Bayes' Theorem:

$$P(s|z, e_1) = \frac{P(s|e_1) P(z|s, e_1)}{P(z|e_1)}$$

$$P(z|e_1) = \sum_s P(s|e_1) P(z|s, e_1).$$

Before performing the experiment, the decision maker can compute the expected value of each action, a_1 and a_2 , using the conditional rather than the prior probabilities of s :

$$u^*(e_1, z, a) = E_{s|z, e_1} u(a, \tilde{s}).$$

The expected value of the choice between a_1 and a_2 , after observing the information represented by z , is the expected value of the best alternative:

$$u^*(e_1, z) = \max_a E_{s|z, e_1} u(a, \tilde{s}).$$

Finally, the expected value of the decision problem with the experiment e_1 is obtained by taking the expectation with respect to its outcomes:

$$u^*(e_1) = E_{z|e_1} \max_a E_{s|z,e_1} u(a, \tilde{s}_1).$$

This is the decision maker's expectation, prior to the experiment, of the utilities to be obtained posterior to the receipt of information. Hence, this technique is referred to as "pre-posterior analysis."

As in the case of perfect information, the expected value of imperfect, or "sample" information (EVSI) is the value of the decision with the information less the value of the decision without it (disregarding cost):

$$EVSI = u^*(e_1) - u^*(e_0).$$

This gives the fair cost of e_1 . If EVSI exceeds the actual cost, $c(e_1)$, then other things being equal, e_1 is worth performing.

2.2.5 Non-negativity of information value.

In a sense, imperfect information about s is perfect information regarding a component of the uncertainty. Uncertainty about s , represented by $P(s)$, is decomposed into uncertainty about the outcome of the experiment, $P(z|e_1)$, and uncertainty about s conditional on those outcomes, $P(s|z,e_1)$. Imperfect information resolves the former. (Note, however, that uncertainty about s in the communication theory sense (Shannon and Weaver, 1949) need not be reduced by the occurrence of z .)

EVSI, like EVPI, is non-negative. The essential point is that

$$\begin{aligned} u^*(e_1) &= E_{z|e_1} \max_a u^*(e_1, \tilde{z}, a) \\ &\geq \max_a E_{z|e_1} u^*(e_1, \tilde{z}, a) \end{aligned}$$

This simply says that taking the maximum within levels of the variable z (perfect information about z) is at least as good as taking the maximum averaged across levels of z (ignorance about z). The remaining steps require some additional assumptions. We see that

$$\begin{aligned} \max_a E_{z|e_1} u^*(e_1, \tilde{z}, a) &= \max_a E_{z|e_1} E_{s|z, e_1} u(a, \tilde{s}) \\ &= \max_a E_{z \& s|e_1} u(a, \tilde{s}) \\ &= \max_a E_{s|e_1} u(a, \tilde{s}) \end{aligned}$$

since $u(a, s)$ is not a function of z . If we further assume that $P(s|e_1) = P(s)$, i.e., that performance of the experiment does not itself alter states of the world,

$$\begin{aligned} \max_a E_{s|e_1} u(a, \tilde{s}) &= \max_a E_s u(a, \tilde{s}) \\ &= u^*(e_0) \end{aligned}$$

Thus,

$$u^*(e_1) \geq u^*(e_0).$$

2.2.6 Non-decisional impact of information. The standard VOI analysis does not allow for the fact that the acquisition of information may, in itself, affect the state of the world. The most obvious example is, perhaps, the situation where, were others

to know that one was obtaining the information, they might alter their own strategies. This is indeed a pressing concern in submarine command and control, for example, where information about target location must often be collected without revealing the presence of one's own ship. Increased accuracy of localization may lead to counterdetection and loss of advantage (Cohen and Brown, 1980). Similar tradeoffs may govern decisions as to the collection of intelligence information in general.

I. H. Lavalley, in his paper, "On Value and Strategic Role of Information in Semi-Normalized Decisions" (1980), examines this "strategic non-independence" between information-gathering and states of the world. Lavalley shows that one may split up the overall value of the information into the strategic value and the pure informational value. The strategic value is the difference in the value of the decision with and without the acquisition of non-used, zero cost information. It may be either positive or negative. The pure informational value is the difference in the value of the decision with and without the decisional use of information already acquired. The pure informational value is thus always non-negative. Lavalley shows that if the decision maker has constant risk aversion (i.e., if his utility functions are either linear or exponential in form) then the overall value of information is simply the sum of the strategic and the pure informational values.

It will be an assumption of our approach that, in general, the utility of information can be additively decomposed into pure information value and other sources of value. The latter includes not only Lavalley's strategic value, but other contributions of an information source or experiment - including, for example, enhanced communication within an organization or increased confidence in previously taken decisions. The utility of information, in terms of its decisional impact, can then be assessed in abstraction from its other sources of value. All contributions to utility can ultimately be combined within an

overall multiattribute evaluation of an information system.

2.2.7 Conditions of positive value. The conditions under which imperfect information has positive (pure informational) value can be illuminated by a formulation in terms of opportunity loss. We assume that in the absence of information concerning z , the decision maker prefers option a_1 . We can now compute the cost of the errors caused by not knowing z . That is, we compare for each value of z , the utility he expects to receive from a_1 with the utility he would have expected had he acted on information about z . The expected value of imperfect information is the expectation of this difference with respect to z :

$$\begin{aligned} \text{EVSI} &= u^*(e_1) - u^*(e_0, a_1) \\ &= E_{z|e_1} [\max_a E_{s|z, e_1} u(a, \tilde{s}) - \max_a E_s u(a, \tilde{s})] \\ &= E_{z|e_1} [\max_a u^*(e_1, z, a) - u^*(e_1, z, a_1)]. \end{aligned}$$

These equations shows that imperfect information will have positive expected value if and only if:

- z and s are not independent ($E_{s|z, e_1} \neq E_s$). Thus, z has an impact on the decision maker's judgments about s .
- Under some outcomes of the experiment, the best action given knowledge of z is not the previously preferred option, so it is possible that $\max_a u^*(e_1, z, a) > u^*(e_1, z, a_1)$. In other words, revised judgments about s will lead to altered decisions.
- The utilities of the previously preferred option and the option indicated by knowledge of z are different: i.e., altered decisions affect payoffs.

2.2.8 Application of VOI to information systems: practicality.

The paradigmatic application of VOI has been to the one-time acquisition of a discrete item of information for a particular purpose: for example, a market survey to decide whether or not to introduce a new product. Even in this kind of application, standard VOI techniques quickly become unwieldy as the number of options, states of the world, and experimental outcomes increase. It is not unheard of for decision trees to be constructed with tens of thousands of nodes.

Information systems, however, especially in C^3I , do not fit within the traditional paradigm, either at the design stage or in use (Figure 1-1). The designer must consider innumerable items of information for inclusion in the actual data base, many of which will be required for a multiplicity of purposes by a variety of users. Moreover, it can be expected that information items within a system will interact strongly in their effects on decision making: in many cases, the presence of one will be useless without the presence of another. As a result, even in a single instance of system use, more than one item will tend to be present in the virtual data base (e.g., on a video display).

Interaction substantially complicates the problem of system evaluation. Consider the task of evaluating a data base consisting of n experiments, $E_n = \{e_1, \dots, e_n\}$, for use in only a single scenario or decision problem. The baseline against which VOI is calculated will depend on the purpose of the evaluation. If the issue is a procurement choice among an information system and other options, we will probably be interested in:

$$EVS I(E_n) = u^*(E_n) - u^*(e_0).$$

That is, we will compare the decision problem with and without the entire data base. But if the issue involves adjustments to an existing system, e.g., the expansion of E_{n-1} by the addition of experiment e_n , we will want:

$$\text{EVSI} (e_n | E_{n-1}) = u^*(E_n) - u^*(E_{n-1}),$$

i.e., we will compare the two data bases directly. Now if the e_i are utility independent, the budget planner or system designer can compute the expected value of the decision separately for each information item, assuming no other items to be present, and then compute the value of the decision with the entire data base by summing over items:

$$u^*(E_n) = \sum_{i=1}^n u^*(e_i).$$

In this case,

$$\text{EVSI}(e_n | E_{n-1}) = u^*(e_n) - u^*(e_0),$$

so when an existing data base is to be modified, only the items in question need be considered.

If, on the other hand, the e_i are utility interdependent, this decomposition fails: the value of adding e_n depends on the items which are already present. Let the set z contain every combination, z , of the outcomes of each of the experiments e_1, \dots, e_n . Then we have

$$u^*(E_n) = E_{z|E_n} \max_{a \in S} E_{s|z, E_n} u(a, s).$$

Thus, a very difficult set of probability assessments is required over $Z \times S$, the joint space of combinations of experimental outcomes and events.

If we now consider the operation of the system across scenarios, the assessment effort is increased many times. (Note in addition that scenarios need not be utility independent: failure in one may increase the value of success in another. Thus, summing EVSI across decision problems may lead to distortion).

Even without interaction, the burden of assessments and computations for system VOI would be prohibitive. With interactions, it is probably impossible.

2.2.9 The problem of specifying structure. As we have seen, VOI methods presuppose a highly structured decision problem: a specified set of options, a specified set of uncertainties bearing on the choice, and a specific source of information which can reduce the uncertainty in specified ways. This requirement raises problems quite apart from the assessment burden it imposes. The set of scenarios in which a C³I system is expected to operate tends not only to be large, but - for a variety of reasons - not very well-specified.

- Technological advances in sensors and weapons will affect options and information sources. Such developments are not fully predictable in advance.
- The options which we perceive as available will evolve as tactical doctrine changes in unforeseen ways.
- The most critical "state of the world" in determining the outcomes of our actions will be the actions of our adversaries. These derive from policies of which we will have only partial knowledge.

In reducing the complexity of VOI, we will decrease the requirement for structure. An incidental bonus is that we may be increasing the scope for information system flexibility.

2.2.10 Standard simplifications of VOI. Methods for simplifying VOI analysis do exist. Some are particular to the value of information, while others represent quite general methods for "pruning" decision trees. Unfortunately, these methods are not entirely successful in resolving the difficulties discussed above. In general, either there is insufficient reduction of complexity or else the analysis is so generic that, in essence, the benefits of the VOI concept are sacrificed.

2.2.10.1 Opportunity loss. A formulation in terms of opportunity loss can sometimes ease the assessments required for VOI analysis. In place of utilities, $u(a,s)$, at the terminal nodes of the decision tree, we can place the error cost or opportunity loss:

$$l(a,s) = u(a_g,s) - u(a,s)$$

where a_g is the option we would choose if we had perfect knowledge that s is the case. We can compute the expected opportunity loss of a decision with experiment e_1 , $l^*(e_1)$, by averaging out and folding back as before, only assuming that the decision maker minimizes expected opportunity loss at each decision node rather than maximizing expected utility (Brown, Kahr, Peterson, 1974; Raiffa and Schlaifer, 1961). The expected value of imperfect information for e_1 is then:

$$\begin{aligned} \text{EVSI}(e_1) &= l^*(e_0) - l^*(e_1) \\ &= [u^*(\text{PI}) - u^*(e_0)] - [u^*(\text{PI}) - u^*(e_1)]. \end{aligned}$$

Opportunity loss is useful when the components of $l(a,s)$, $u(a_g,s)$ and $u(a,s)$, are additively decomposable into: (i) a common factor which is difficult to assess, and (ii) a critically varying factor. For example, some aspects of the utility of a decision problem may be constant whether or not the optimal action is selected, e.g., the cost of deploying a platform may not depend on the tactics adopted. This common factor drops out of the analysis and need not be considered when the difference, $l(a,s)$, is assessed directly.

Analysis in terms of opportunity loss omits no relevant information from the assessment of EVSI. Clearly, however, the reduction in the assessment task is quite marginal in terms of the total assessments still required.

2.2.10.2 Perfect information. The expected value of perfect information can be used to set an upper bound on the amount

the decision maker should be willing to pay for information. The computation of EVPI is considerably simpler than for EVSI, whether or not in terms of opportunity loss. The experimental outcomes need not be enumerated, and a probability measure need be assigned only to S rather than to $Z \times S$. Moreover, a rough assessment of EVSI can be based on EVPI by the process of "anchoring and adjustment." The evaluator need only assess the proportion P_1 of the value of perfect information which he believes will be afforded by the imperfect information e_1 . Then

$$EVSI(e_1) = P_1 \cdot EVPI$$

The evaluator will probably place more confidence in the adjustment process if he has actually calculated both EVSI and EVPI for a few information items. Nonetheless, it is far from clear that the assessment of P will be sufficiently refined to distinguish among many proposed alternative data bases. In comparing e_1 and e_2 , for example, the relevant comparison is between P_1 and P_2 , EVPI dropping out as a common term. But this boils down simply to a holistic assessment of VOI for e_1 and e_2 , with no explicit analysis of their differential impacts on decision making.

2.2.10.3 Scenario sampling. An exhaustive elaboration of scenarios is fortunately not required for the evaluation of information systems. Considerable decision-analytic experience has been gained in the past few years in the generation of sets of manageable scenarios which, nevertheless, constitute adequate bases for evaluation of systems.

O'Conner and Edwards (in press) have proposed some general concepts and methods for generating a satisfactory set of scenarios. Four criteria can be identified from their paper. Scenarios should be:

- (1) realistic;

- (2) relatively probable;
- (3) representative;
- (4) maximally discriminating among systems.

Criterion (1) can be satisfied by adding detail that has approximately the same effect on the evaluation of all alternatives. Criterion (2) requires not that scenarios be probable in a literal sense, but that they be probable relative to other scenarios specified at a comparable level of detail. Criterion (3) requires that the selection of scenarios be an appropriate sampling from the total "scenario space." Once the appropriate dimensions of this space are defined, achievement of criterion (3) can usually be assessed by expert judgment. The major difficulty is in defining these dimensions so as to satisfy criterion (4) as well.

Once appropriate scenarios are selected, system utility can be estimated as a weighted average of the utility within each scenario. Weights represent the probability of a scenario relative to the probability of the selected scenario set, rather than its probability as such (which will be very small).

These concepts were developed in the context of applying multi-attribute utility models to the evaluation of complex military systems. Their relevance to present concerns, however, is twofold. First, of course, they may be applied straightforwardly to the evaluation of C³I systems. More important, however, is a second -- and novel -- application of these concepts to the determination of the expected value of a set of information. This latter use would, in essence, substitute sets of information in place of the alternative radios (Chinnis, Kelly, Minckler, and O'Conner, 1975), radars (Barclay, Chinnis, and Minckler, 1975), or other equipment being evaluated.

These techniques, however, do not reduce in any way the assessment effort required within a scenario; and the total number of scenarios which needs to be considered may still be quite large.

2.3 Modifications in VOI Techniques

The problem of simplifying VOI analysis so that it becomes manageable for information systems remains unsolved. In this section, we turn to some more fundamental problems in the VOI approach. Modifications of standard techniques which have been proposed to deal with these problems may lead us some distance toward the desired simplifications.

2.3.1 The consistency condition. A very strong assumption, implicit within the standard decision analytic treatment of VOI, concerns the behavior of the decision maker when he arrives at the subsequent decision node (e.g., the choice between a_1 and a_2). Not only is it assumed that the structure of the decision problem is fully specified, but it is assumed that the decision maker acts rationally upon that structure. The "consistency condition" says that the portion of the decision tree following the decision node will adequately represent the decision maker's view of his problem when he gets to it. This means:

(1) He will compute expected utilities for each option using the modeled conditional probabilities and utilities.

(2) He will then maximize expected utility.

(Alternatively, we need only require that he will act as though (1) and (2) were true.)

The traditional analysis thus takes a normative view of subsequent as well as initial decisions. Its output is not simply a recommendation concerning information purchase, but a "decision rule" prescribing subsequent choices contingent on outcomes of the experiment.

This normative approach, and the consistency condition underlying it, are closely related to the non-negativity of VOI.

It must be remembered that VOI is the prior expectation of the value of information, before receiving it. If the decision maker feels that the information to be provided will be misleading, he should also anticipate that he will take that feeling into account when he comes to act on the information. (It is reflected in the conditional probabilities for states of the world which he assesses now and will act on then.) By incorporating his misgivings into the model, he dispells them. In essence, the consistency condition states that, at any given time, he should not be surprised about (have failed to anticipate) his feelings or opinions at that time.

Nevertheless, the normative approach to subsequent decisions, reflected in the consistency condition, may conflict with the goal of evaluating information decisions. For this purpose, what is needed is a prediction of what the decision maker will do with the information. In effect, the consistency condition takes prescription as description: he will optimize with respect to the information available at the time of the analysis.

2.3.2 Acts as events. Rex Brown, in his paper, "Heresy in Decision Analysis: Modeling Subsequent Acts Without Rollback" (1975), advances the idea that, in certain situations, subsequent acts should be modeled by probability nodes rather than decision nodes. Thus, rather than constraining the decision maker to maximize expected utility for decisions some time in the future, we suppose that there is uncertainty in the way in which he will act.

2.3.2.1 Failures of consistency. The circumstances in which this proposal is appropriate are just those in which the consistency condition fails to hold:

(1) Such a concept is of value if one believes that the decision tree does not fully model the events up to the time of the decision, i.e., the information that may then be available

to the decision maker. If this possibility is admitted, then the concept of maximizing expected utility for subsequent acts, by rolling back the decision tree, becomes invalid. Even if the decision maker in fact maximizes his expected utility in the subsequent decision, the option he selects need not be the one which maximizes expected utility conditional only on the modeled information.

(2) Similarly, the decision maker, now, may feel that in the future, he might change certain aspects of his analysis of the problem, including his probability assessments or utility function. In principle, it is possible to model any such set of changes as event forks in the decision tree prior to the subsequent decision node. But it is clearly out of the question to model explicitly all possible probability assessments and utility functions (cf., Brown, 1975). Not only would probability distributions have to be assigned to all the possible changes, but--in order to predict the decision maker's choice by rolling back the tree--probabilities subsequent to the decision node would have to be made conditional on these changes. This would require a set of second-order assessments on the diagnosticity of the new probabilities or utilities, regarded as information, for the "true" or "authentic" values (cf., Nickerson and Boyd, 1980, and Tani, 1975, discussed in Section 2.3.2).

These observations suggest an in principle limitation to the modeling of this type of change, in one's basic view of the problem, within the traditional VOI paradigm. There is nothing logically to prevent the decision maker from changing these second-order conditional assessments as well during the time before the subsequent decision. But in order to model this change, he would need a new, higher-order model--and so on, into an infinite regress. Far-fetched as this kind of example seems, it does represent a fundamental formal limitation on the power of traditional techniques to describe all decision problems.

(3) Some instances in which the consistency condition is violated do not seem amenable to fuller modeling in terms of events at all. The decision maker may predict that on account of fatigue, time pressure, limitations of memory or cognitive capacity, he will fail to process certain information (e.g., a danger signal) appropriately. This could be true even though in the course of the VOI analysis, he (perhaps with the help of a decision analyst) has already worked out how he should process it. It seems quite implausible to regard factors like fatigue as "information events." In these cases, the decision maker simply lapses from his own normative standard.

(4) Another case which does not seem amenable to traditional methods is that in which the decision maker adopts a new "decision rule." Perhaps his behavior, rather than maximizing expected utility, can be expected to reflect some of the "heuristics" described in recent psychological literature (Kahneman and Tversky, 1977; Slovic, Lichtenstein, 1971).

One can always, of course, so define the decision maker's utility function that the "act" of obeying a heuristic turns out to maximize expected utility. In this case, one attributes to the decision maker an affection for the heuristic as such. But this rather trivializes the notion of optimal behavior and, further, is likely to contradict the decision maker's own, more reflective normative views. In addition, there is a near contradiction in supposing that the decision maker uses maximization of expected utility (on a second-order level) to conclude that he should not use it (on a first-order level).

The plausibility of the consistency condition may be particularly doubtful in the case of information systems, where the system designer and the decision maker are not, typically, the same person. The likelihood that decision maker and system

designer use different decision rules, or that the system designer will fail to consider information events available to the system user, is correspondingly greater.

2.3.2.2 Sufficient modeling. Even if it were possible, in principle, to model all these problematic changes as "information events," it is not logically necessary to do so. Modeling acts as events represents a formally sufficient degree of modeling (Brown, 1978), even when information events are omitted from the analysis or when acts are influenced by other sorts of causes. Events may be implicitly "integrated out," without affecting the results of the analysis, as long as they do not precede a decision node. And there should be no greater difficulty in predicting one's own behavior, as an uncertain event, than in predicting environmental events which, similarly, are subject to complex and inexplicit causal influences

To be sure, the acts as events model does have shortcomings. Most obviously, using this idea will to some degree lessen the prescriptive power of the decision analysis paradigm, for the decision maker will not get so much guidance in his future decisions.

More importantly, perhaps, little guidance has been provided, thus far, on how the required assessments are to be obtained. Two sorts of novel judgments are needed in this approach:

- for each option, the probability that the decision maker will adopt it, conditional on the modeled information; and
- for each option, its expected utility, conditional on the assumption that the decision maker has adopted it.

These assessments essentially reverse the order of precedence in the standard technique. There, the probability of a state of the world, conditional on the outcome of an experiment, was independent of the action adopted:

$$P(s|e,z,a) = P(s|e,z),$$

since the action was assumed to be selected solely on the basis of an inference about s suggested by the experimental outcome. When, however, it is assumed that action is based on unmodeled information, the assessment of the probability of s , or the probability distribution on terminal utility u , must be adjusted to take this hypothesized information into account. If the act is selected, we assume it was for a reason. It is not clear, however, what the nature of this adjustment should be.

Despite these difficulties, it is clear that the acts as events model holds promise for our considerations of the value of information stored in a data base. The expected benefits include the simplification due to not modeling all information items and the entire decision process, as well as the increased validity due to accomodating violations of the consistency condition. Moreover, as we shall see, methods to facilitate the assessment task are available.

2.3.3 Value of analysis. Among the items which may potentially be included in an information system are programs for inference and decision aiding. Such programs, which are frequently interactive, may help the decision maker structure the problem, by identifying options and major uncertainties bearing on the outcome, and may help him arrive at a solution by applying formal tools of statistics and decision theory (e.g., Barnes, 1980; Cohen and Brown, 1980). Inputs may be either subjective judgments or objective data, including other items in the data base.

The evaluation of such programs raises a special set of problems in the application of decision analytic technique:

- Standard VOI techniques assume that the situation is already modeled before the analysis can be evaluated. Were such indeed the case, there would of course be no value in performing the analysis again.
- Decision analysis decomposes initially very complicated and difficult assessments into simpler ones. If the result of the analysis were always consistent with initial judgment, there would be no point in doing the analysis. We thus view the decision maker as an inconsistent probability assessor. It seems odd, then, to base a theory of the value of analysis on the consistency condition, as embodied in VOI.

Using decision analysis can be very expensive. It is therefore not surprising that in recent years potential clients have wanted some indication of how much the analysis would be worth, before they contract for it. However, only a few papers have been published upon the subject of evaluating decision analysis. These appear relevant, as well, to the issue of including automated decision analytic aids--in addition to aids of other types--within information systems. In this section we shall review the three major papers appearing in the literature.

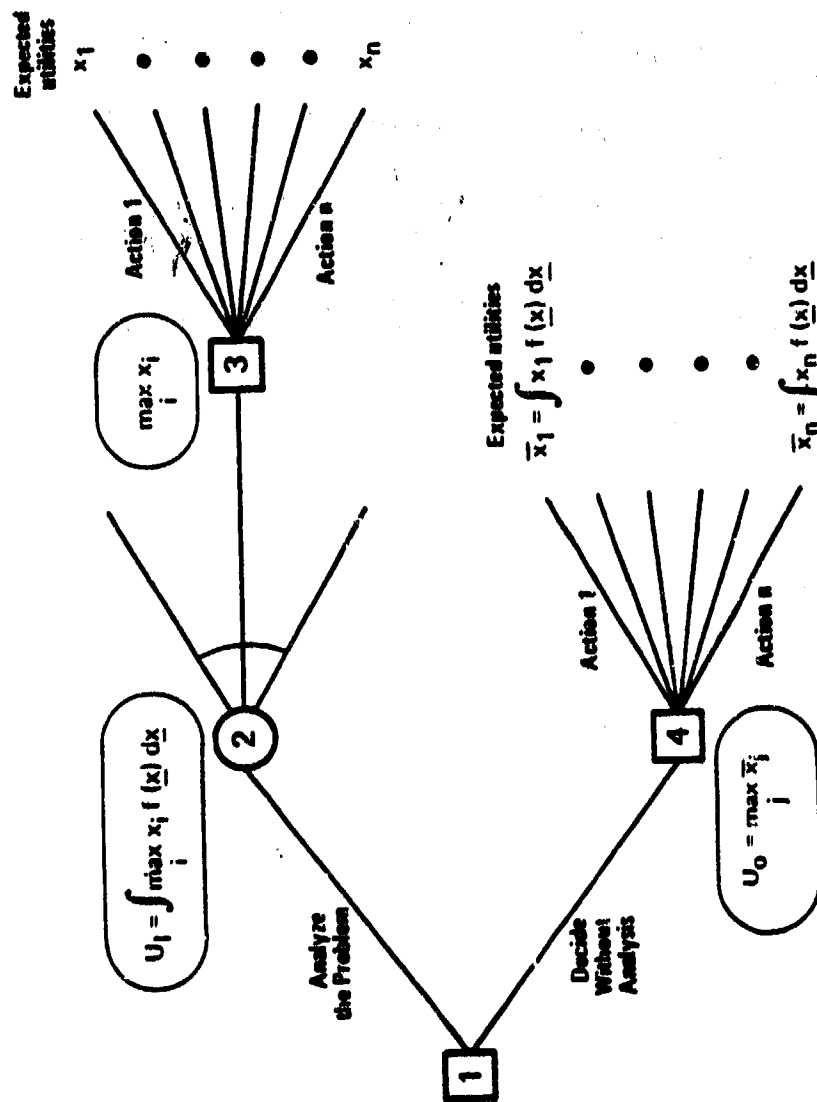
2.3.3.1 Watson and Brown. The paper by Watson and Brown, entitled "The Valuation of Decision Analysis" (1975), uses a decision analytic model in the evaluation. However, rather than require a full model of the problem, they use the simple tree of Figure 2-3. Uncertainty is encapsulated within the x_i , which are the expected utilities for actions which will arise from "perfect analysis." In evaluating such perfect analysis,

the calculation is precisely the same as with value of perfect information, but information concerns the x_1 . Because this is simply value of perfect information, it is easily shown that such value will always be positive.

Watson and Brown similarly model the value of imperfect analysis by assuming that the imperfect analysis will alter our prior distribution concerning the results of a potential perfect analysis. Thus a node is introduced before the decision node which describes the outcome of the imperfect analysis, and another uncertainty node appears after the decision node which describes the uncertainty concerning the perfect analysis, given the imperfect analysis results. However, Watson and Brown do not assume the consistency condition and thus are able to provide an example where the value of imperfect analysis is negative.

Rather than assume the consistency condition the authors assume that the decision maker will follow the recommendations of the imperfect analysis, whatever his prior beliefs may have been concerning its validity. They thus assume a version of "acts as events" in which the probability of an act is equal to the likelihood that it will be recommended by imperfect analysis. Negative expected value arises if the decision maker believes beforehand that the imperfect analysis will be sufficiently poorly done as to worsen his position.

If those prior beliefs were incorporated into the analysis, we would find that the prior expectation of the results of the perfect analysis, given the imperfect analysis, would be precisely the results of the imperfect analysis. After all, if, after an imperfect analysis, the expected value of the expected utility of an action that would be obtained by a perfect analysis, differed from the expected utility obtained by the imperfect analysis, the imperfect analysis could be instantly improved by changing its results so that the equality did hold.



- Notes:
- (i) The expressions ringed at the nodes of this tree are the expected utilities there.
 - (ii) Node 2 represents the analysis which produces \underline{x} , the vector of expected utilities associated with the actions. $f(\underline{x})$ is the decision maker's density function for \underline{x} .

Figure 2-3

DECISION TREE FOR THE VALUATION OF ANALYSIS

Incorporating prior beliefs in this way is equivalent to the consistency condition; selecting the act with the highest expected value according to the imperfect analysis is, in this case, the same as rolling back the decision tree. In such a case, it is easy to prove that the expected value of analysis will be positive.

This version of the consistency condition seems sensible. It is plausible that an (imperfect) analysis might in fact lead one astray. But it is less plausible that one would expect this to happen beforehand, yet also expect to undertake the analysis and abide by its results.

The real issue of consistency in this context is not whether the decision maker will roll back the tree (or modify the imperfect analysis) based on his conditional expectation of the results of perfect analysis. The issue rather is whether he can meaningfully formulate such an expectation in the first place. Joint probability distributions need to be assessed over the values of imperfect and perfect analysis. Since the distribution over the results of the imperfect analysis assumes a good understanding of what that analysis will be, it seems unfair that the decision maker be required to make such assessments. Yet it is not in the spirit of decision analysis for the analyst to describe how good he believes his analysis will be, or to prescribe the degree of confidence in it to be felt by the decision maker.

The concept of perfect analysis may itself not be meaningful to the decision maker. Not only can it not in practice be carried out, but the existence of a unique set of psychological probabilities to be uncovered by such an analysis is doubtful.

Despite these difficulties, it may be that inconsistent assessments (which motivate the use of decision analysis) cannot be modeled without some concept of perfection, to provide a point of comparison against which the obtained assessments can be evaluated. Such a

concept is indeed common to all approaches to this, and similar problems. The calculation of the value of analysis may be sufficiently insensitive to the precise nature of the second-order distributions, to justify the use of these concepts at least in a heuristic sense.

The Watson and Brown approach assumes that the available options to the decision maker have already been modeled. However, in many situations this is a very major part of the analysis. Once this has been carried out, it may well be easier to perform a very quick and rough decision analysis on the problem, carry out a sensitivity analysis on the parameters and use this information to help the decision maker decide, and use this further analysis is required. Thus, difficult second-order assessments need not be undertaken if options are specified.

2.3.3.2 Tani. Steven Tani, in his paper, "A Perspective on Modeling in Decision Analysis" (1978), views the aim of decision analysis as being to help us obtain a probability distribution over an outcome variable (e.g., profits) that is "better" than the one which we can assess directly. In order to define "better", he brings in the concept of an "authentic probability." An authentic probability is one which most accurately describes all our beliefs about the event in question. He then contrasts these authentic probabilities with probabilities that can actually be assessed, which he calls "operative." The "goodness" of an operative probability is then defined in terms of its closeness to the authentic probability.

Tani builds a probability distribution describing where the authentic probability may in fact lie given the elicited operative probability. We view the authentic probability as that probability we would provide after an infinite time for thought and introspection. We may view this probability distribution over the authentic probability as our uncertainty, after

finite time for thought and introspection, as to what we might think after infinite time. Its expected value is the operative probability itself (a version of the consistency condition similar to that discussed by Watson and Brown); and the variance of this distribution measures closeness to the authentic probability.

The value of modeling is simply the value of reducing our uncertainty concerning the authentic probability. Standard VOI analyses may then be carried out on these probability distributions and the value of modeling can be calculated.

Tani's concept of an authentic probability provides some elucidation of Brown and Watson's notion of "perfect analysis." It is also similar to work in reconciling incoherent probability assessments by Lindley, Tversky, and Brown (1979). They use the idea that operative probabilities are noisy measurements of the authentic probabilities, and may therefore be in error. The problem, of course, is that the required assessments may still be quite difficult. There is no guarantee that the second-order probability distributions will themselves be authentic! A second problem is that, again, the calculations may only be carried out if the majority of the modeling has already been performed.

2.3.3.3 Nickerson and Boyd. Nickerson and Boyd, in their paper, "The Use and Value of Models in Decision Analysis" (1980), take the view that the modeling in decision analysis should be viewed by the decision maker as simply another piece of information to be incorporated into his overall system of beliefs. Thus, rather than expecting the decision maker to follow the recommendations of the analyst unquestioningly, Nickerson and Boyd suggest that the decision maker should use Bayesian updating on his beliefs prior to the analysis and use his updated beliefs, independently, to choose his preferred option posterior to the analysis.

The model that Nickerson and Boyd build for evaluating a decision analysis supposes that there is a "true" equivalent reference value for each option: the value that will actually result if that option is selected. Then the authors assume that, at any stage, the decision maker can build a probability distribution over what that true value might in fact be. Then "perfect" modeling is defined as a modeling effort that yields a probability mass function with all mass concentration at the actual equivalent reference value for each alternative. Thus, perfect modeling provides perfect information about the equivalent reference value. This is therefore a much more stringent requirement for perfect modeling than the ones used in the other papers, where a perfect model is simply one which fully encapsulates our current subjective beliefs about the expected utility of each action.

Nickerson and Boyd are then able to use the standard VOI concepts to calculate the expected value of perfect modeling, as the expected value of perfect information about the equivalent reference values. Similarly the expected value of imperfect modeling is simply the value of imperfect information about those equivalent reference values, under the assumption that the decision maker will incorporate that information optimally into his belief structure. It can then be easily shown that the expected value of modeling must always be non-negative.

A major difficulty with this approach is the assumption of the existence of "true" equivalent reference values, to serve as points of comparison for evaluating obtained assessments. Such values are clearly not objective, in the sense that the yearly amount of rainfall in Minneapolis is objective. They incorporate such subjective factors as utility functions and tradeoffs between different attributes. In order to determine a "true" equivalent reference value, therefore, one cannot simply predict, and then observe, the value which "occurs." Some digging within the decision maker's psychological field is required, to

ensure that all his relevant beliefs are incorporated. But this is essentially the idea of "perfect analysis" or "authentic probabilities."

On this approach, as on the previous ones, the assessments required with imperfect modeling, in order to incorporate the imperfect information into the decision maker's beliefs, are very difficult. The method also suffers, of course, from the usual difficulty that the options are assumed predefined.

However, there do appear to be many situations in which these concepts may prove applicable. The idea of using decision analysis only as a source of information to the decision maker seems more valid than the more usual assumption that the recommendation of the decision analysis is what the decision maker should choose to do. A very similar procedure was developed by members of DSC to help evaluate potential information in a forest planning situation (DSC Report, 1980). The notion that decision aids are information seems particularly appropriate in the area of command and control, where it underlines the fact that such aids support, but do not replace, the decision maker.

3.0 TECHNIQUES FOR INFORMATION SYSTEM DESIGN

In this chapter we sketch out, on a conceptual and algorithmic level, some ideas for assessing information value in the context of system design. In large part, these ideas are as yet untested.

We observed in Chapter Two that standard VOI techniques, while properly focusing attention on the decisional impact of information, require a highly structured context in order to be applied. They presume a specified information source or experiment, specified experimental outcomes, specified uncertainties about states of the world, and specified options. Such analyses quickly become intractable when applied to complex, multi-purpose systems, expected to operate in a variety of environments, some of which are quite ill-defined. Moreover, they lack the flexibility to handle non-optimal responses to information. And, finally, they cannot be employed to evaluate aids which support inferential and decision making processes.

Our objective is to devise methods which are simple and flexible, yet incorporate in a realistic manner the impact of information on decision making. Clearly, different degrees of specification and structure will be appropriate for different evaluative purposes. In each case, specifications which limit the tractability or validity of the analysis at hand must be omitted. Our strategy, however, is to find, within this constraint, a degree of modeling that is not only formally sufficient to express the value of information, but which includes explicit reference to decisional impact.

We will begin with a relatively highly structured technique and derive new ones by progressive abstraction. The first technique to be considered, however, is, itself, a generalization over standard VOI procedures. The models to be described reflect the

following sequence of increasing abstraction:

- (0) standard VOI;
- (1) unspecified experimental outcomes;
- (2) non-optimal behavior;
- (3) unspecified states of the world;
- (4) unspecified options;
- (5) multiattribute approximation.

3.1 Unspecified Experimental Outcomes

In a number of circumstances traditional VOI techniques are inapplicable on account of the character of the data set being evaluated. Those techniques require:

- (1) that the information take the form of "experiments", with variable outcomes not known in advance; and
- (2) that the possible outcomes of an experiment be spelled out and probabilities assigned to them.

Often, however, it is either inconvenient or impossible to satisfy these conditions. In particular, the impact of a given item of information may depend on the context of other items in the data base. Spelling out all possible combinations of outcomes of all experiments, assessing their probabilities, and assessing probabilities for states of the world conditional on such combinations, may present insuperable difficulties (see Section 2.2.7).

Moreover, in some cases the information to be evaluated will take the form not of experiments, but of facts, from physics, history, or recent intelligence. Facts make trouble for traditional preposterior analysis since, unlike experiments, they have but a single "outcome." Facts are known to the system designer, but are not - if inclusion in a data base is useful - present in the "working memory" of the system user. In this case it is obvious that predicting user choices by rolling back the system designer's model of the subsequent decision will be inappropriate. Since the fact is already known to the system designer, conditionalizing on it will not alter the probabilities he assigns to states of the world. It is the user whose posterior probabilities (and actions) may be affected by inclusion of the fact in the Actual Data Base.

3.2 Acts as Events in System Evaluation

These difficulties can be handled by treating acts as events (Brown, 1975). Information, to the extent that it has value, may change a decision maker's choice of action (Section 2.2). Thus, if experimental outcomes occurring prior to a decision are not explicitly modeled, the prediction of action based on rolling back the decision tree is invalid (Section 2.3.2.1). However, an alternative approach allows such outcomes to be "integrated out." This involves treating subsequent decisions as uncertain events and assessing probabilities for states of the world and utilities conditional on the action selected (Section 2.3.2.2).

The acts as events model applies to the evaluation of a single item in a data base, to the evaluation of groups of items, and to the evaluation of facts.

(1) We first consider the value of adding information item e_n to a data base, E_{n-1} , consisting of items (e_1, \dots, e_{n-1}) . The

value of e_n may well depend on the inclusion of some or all of E_{n-1} (Section 2.2.7). If such interactions exist, the computation of this value,

$$EVSI(e_n | E_{n-1}) = u^*(E_n) - u^*(E_{n-1}),$$

according to the standard VOI analysis, requires an enormously complicated decision tree in which the possible outcomes $z^{(i)}$ of each experiment e_i are modelled (Figure 3-1). If, however, we treat the decision node for a as a chance node, we may validly omit reference to the outcomes of experiments.

In figure 3-2, we choose to model only the outcomes of e_n , i.e., the experiment under consideration. Uncertainty concerning the outcomes of other experiments must be implicitly considered in assigning probabilities to acts, a , events, s , and (if S does not exhaust the relevant states of the world) to the distribution on terminal utility, u . If the evaluation can consistently produce such assessments, the value of $EVSI(e_n | E_{n-1})$ calculated from Figure 3-2 should be the same as from Figure 3-1.

(2) The acts as events model enables us to group experiments for the purposes of assessment. Consider the evaluation of a large set of information items, E_n , as the potential contents of a data base. If the elements of E_n interact in their impact on decisions, standard VOI once again requires a complex tree, as in Figure 3-3, to evaluate:

$$EVSI(E_n) = u^*(E_n) - u^*(e_o).$$

By treating acts as events, we can achieve any desired degree of simplification. All experimental outcomes may be omitted, or (as in the lower branch of Figure 3-2) we may explicitly model outcomes for selected elements in E_n .




Figure 3-1

Standard VOI for a Single Experiment




Figure 3-2

Acts as Events VOI for a Single Experiment




Figure 3-3

Standard VOI for a Group of Experiments

(3) The inclusion of facts in a data base could never, on standard VOI analysis, be justified. Mere inclusion of an already known fact in a data base will not change probabilities for states of the world, hence, will not change subsequent decisions. (And we do not wish to compare inclusion of a fact with a hypothetical state of the world in which the fact was not true!) By treating acts as events, however, we can distinguish the effect of including a fact on the user's decision making, from its role in the system designer's assessment of probabilities for states of the world. Since acts are not predicted by rolling back the tree, the system designer's probabilities for states of the world remain constant while the user's actions change.

The value of knowing a fact will, for most C^2 applications, depend on the experiments which are also included in the data base. Thus, a fact, f_1 , will typically be evaluated in a context of experiments (and other facts) E_n , in accordance with the methods of (1) and (2) above. In Figure 3-4, experimental outcomes for one experiment, e_j , have been explicitly modeled and others integrated out.

Note that the probability assignment for states of the world is independent of the inclusion of f_1 in the data base:

$$P(s|E_n + f_1, z^{(j)}, a) = P(s|E_n, z^{(j)}, a).$$

That is, in either case the system designer uses all his factual knowledge, including f_1 , to evaluate the conditional probability of s . On the other hand, the probability of an action is not independent of the inclusion of f_1 :

$$P(a|E_n + f_1, z^{(j)}) \neq P(a|E_n, z^{(j)}),$$

assuming that the decision maker does not already have a working knowledge of f_1 .




Figure 3-4

Acts as Events VOI for Addition of a Fact to Data Base

3.3 Assessment heuristics. The acts as events methodology vastly reduces the number of judgments required in order to evaluate experiments. However, the assessments which remain appear quite difficult. A large amount of implicit information must be integrated, within the mind of the assessor, in order to arrive at estimates of probabilities and expected utilities for actions. The probability that an action will be chosen must reflect the expected impact of omitted experimental outcomes, and the expected utility of the action must be based on the assumption that the act was chosen on account of such impact.

3.3.1 Auxiliary decision tree. One approach is to assess these quantities directly. However, a more natural procedure would be to start with the assessments in a traditional, but abbreviated (hence invalid) decision tree, and then derive the required assessments by adjustment. Consider again the value of adding a single experiment, e_n , to the data base, E_{n-1} (Figure 3-2). For this problem we construct the auxiliary tree in Figure 3-5. This, of course, is the diagram one would use in a standard preposterior analysis which, incorrectly, fails to model e_n 's information context, E_{n-1} . Rolling back this tree, we compute the expected utility for each option, conditional on incomplete information,

$$u_m^*(E_n, z^{(n)}, a) = E_{s|e_n, z^{(n)}} u(a, \tilde{s})$$

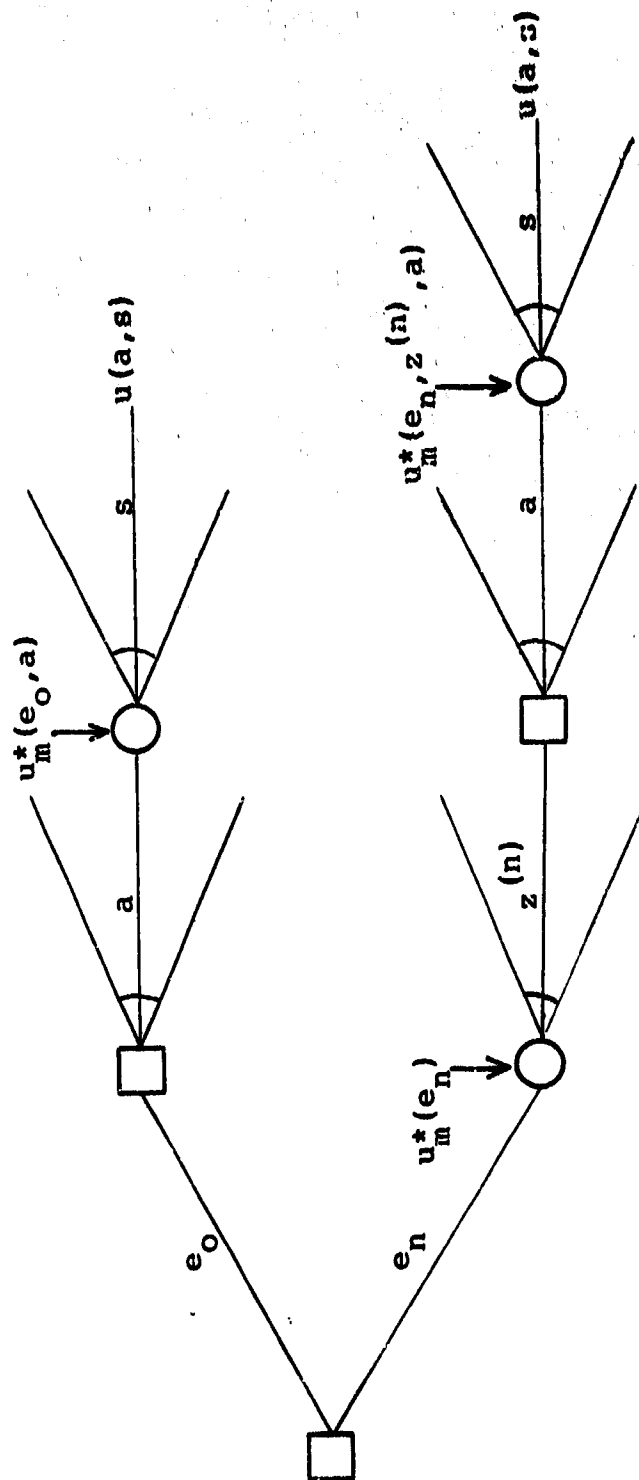
and similarly,

$$u_m^*(e_0, a) = E_{s|e_0} u(a, \tilde{s}),$$

where subscript m indicates that these values derive from a auxiliary model. The assessments for utility and for $P(s|e, z)$ are straightforward here, since acts are not regarded as events upon which they must be made conditional.

Figure 3-5

Auxiliary Tree for Acts as Events Model in Figure 3-2



It remains now to discuss methods for deriving the required judgments: act probabilities and expected utilities conditional on complete information.

3.3.2 Assessment of act probabilities. If, contrary to assumption, Figure 3-5 modeled all the relevant information available to the decision maker, we would expect him to choose, after observing $z^{(n)}$, the action which maximizes $u_m^*(e_n, z^{(n)}, a)$. The problem, of course, is that, because we have integrated over information events, there is uncertainty as to the actual expected utility, $u^*(E_n, z, a)$, which the decision maker will assign to option a after observing all the experimental outcomes. $z^{(n)}$ is a partition of the total space, $z = z^{(1)} \times z^{(2)} \dots \times z^{(n)}$, of experimental outcomes, having been obtained by marginalizing over outcomes of all other experiments. Thus, $u_m^*(e_n, z^{(n)}, a)$ is merely the expectation, with respect to the ignored experimental outcomes, of the random variable, $u^*(E_n, \tilde{z}, a)$. The probability that the decision maker will select option a_i is the probability that $u^*(E_n, \tilde{z}, a_i)$ exceeds $u^*(E_n, \tilde{z}, a_j)$ for all $i \neq j$.

A well-known device for approximating such probabilities is based on the choice axiom described by Luce (1959, 1963, 1977). Abbreviating $u_m^*(e_n, z^{(n)}, a_i)$ as $u_{m,i}^*$, we let:

$$P(a_i | E_n, z_n) = \frac{(u_{m,i}^*)^c}{\sum_j (u_{m,j}^*)^c}$$

The exponent, c , is a measure of the accuracy with which the expected values, $u_{m,i}^*$, represent the distributions of $u^*(E_n, \tilde{z}, a)$. It thus reflects the system designer's confidence in the completeness of his model of information events. (Yovits, Rose, and Abilock (1969-78) use this equation, with an analogous interpretation of c , for a quite different, normative purpose. In the context of signal detection (Luce, 1963), the exponent reflects the number of independent observations of a stimulus.)

Using the choice equation, probabilities for acts can be assigned on the basis (i) of their expected utilities conditional on incomplete information and (ii) a single additional assessment, c. The plausible range for c is between zero and infinity. When the modeled information is complete, c goes to infinity, and the choice equation predicts with a probability of one selection of the act with highest expected utility. On the other hand, when the number of unmodeled events and the swings in utility which they can generate are large, there is no reason to suppose that the judgments of the decision maker will correspond to the model. In this case, c equals zero, and choice probabilities, from the point of view of the model, are random:

$$P(a|E_n, z^{(n)}) = 1/q \text{ for } q \text{ alternative actions.}$$

The intermediate case in which $c=1$ is of particular interest, since it represents the case in which response probabilities are proportional to utility (cf., Luce, 1959).

A nice property of the choice equation is that, for $0 < c < \infty$, the order of action probabilities corresponds to the order of expected utilities conditional on incomplete information. Use of the equation thus ensures satisfaction of a desideratum for the assessment of action probabilities suggested by Brown (1975).

Note that by introducing terms for "response bias" (Luce, 1963), we obtain a more general form:

$$P(a_i|E_n, z_n) = \frac{b_i (u_{m,i}^*)^c}{\sum_j b_j (u_{m,j}^*)^c}$$

The b_i reflect any feelings the system designer may have about the inclination of the decision maker to adopt specific acts, independent of events (whether modeled or unmodeled). Thus,

when c equals zero, actions are unpredictable in terms of events, and:

$$P(a_i | E_n, z^{(n)}) = \frac{b_i}{\sum_j b_j}$$

3.3.3 Assessing c . A convenient procedure for assessing c is suggested by the following formula, in which we write P_i for $P(a_i | E_n, z^{(n)})$:

$$\log \frac{P_i}{P_j} = c \log \frac{u_{m,i}^*}{u_{m,j}^*}$$

c is, therefore, the slope in logarithmic coordinates of the plot of response probability ratios to ratios of expected utilities.

This relationship can be exploited in two different ways:

- Ratios of response probabilities can be assessed for several hypothetical ratios of expected utilities. A straight line may then be fit in log-log coordinates to the obtained points, and the slope taken as an estimate of c .
- The slope may be assessed directly by adjusting a straight line in these coordinates. The implications of any particular adjustment for the relation between the two sets of ratios may be read off.

Both methods assume that the assessor will have reasonably consistent intuitions that, for example, a response twice as high in expected utility is three times as likely to be chosen, etc.

It is worth noting that the arctan transformation of c , corresponding to the angle of the slope, transforms the range of c to a bounded interval, with $c = 1$ as the midpoint. This may be a reasonable scale to use, therefore, if it is desirable to get a quick numerical estimate of "model completeness."

3.3.4 Assessment of expected utilities. Unfortunately, we know of no procedure of comparable simplicity for deriving expected utilities for actions, conditioned on the action having been taken, from the modeled expected utilities. We do know, of course, that the adjusted expected utility, $u^*(e_n, z^{(n)}, a_i)$, will be greater than the modeled one, $u_{m,i}^*$, if unmodeled events are thought to affect the utility of a_i . If a_i is chosen, we assume that the unmodeled uncertainties came out in favor of a_i , and its expected utility is greater than the average with respect to those uncertain events.

Ideally, a formula could be provided which uses the parameter c and the $u_{m,i}^*$ to give $u^*(E_n, z^{(n)}, a)$. In the absence of such a formula, however, some rules of thumb can be offered. These rules require a rough assessment of the effect of unmodeled events on the expected utilities of the options, i.e., the uncertainty concerning $u^*(e_n, \tilde{z}, a_i)$, given $z^{(n)}$.

- If uncertainty about expected utility $u^*(E_n, \tilde{z}, a)$ for given experimental outcome $z^{(n)}$ is approximately the same for all options a , the order of the adjusted expected utilities, $u^*(E_n, z^{(n)}, a)$, should be the same as the order of the modeled expected utilities, $u_{m,i}^*$.
- The larger the effect of unmodeled events on $u^*(E_n, \tilde{z}, a_i)$, the greater the adjustment for that act: i.e., $u^*(E_n, z^{(n)}, a_i)$ exceeds $u_{m,i}^*$ by a greater amount.
- Other things being equal, the lower the modeled expected utility of a_i , the more it must increase to be the maximum, and the greater the adjustment. Thus, the spread of the $u^*(E_n, z^{(n)}, a)$ is less than the spread of the u_m^* .
- When expected utilities for acts are negatively correlated, i.e., events that help one act hurt another - the adjustment needs to be less.

These guidelines are based on the behavior of the expected value of a random variable, conditional on its being the maximum of a set of random variables.

A lower bound on the required adjustment can be quickly derived if one is prepared to assume that the distribution of $u^*(E_n, \tilde{z}, a)$ for given experimental outcome $z^{(n)}$, is normally distributed with mean $u^*(e_n, z^{(n)}, a)$ and constant variance σ^2 for all a . Let a_1 be the act which maximizes $u^*(e_n, z^{(n)}, a)$ in the auxiliary model

(Figure 3-4). We now make the further assumption that the correlation between $u^*(E_n, \tilde{z}, a_1)$ and $u^*(E_n, \tilde{z}, a_i)$ for all other acts a_i is -1 . This means that any gain in utility for act a_i , due to unmodeled events, is offset exactly by a loss of utility for a_1 . It follows that the lowest expected value for a_i at which it could be preferred to a_1 is midway between the two modeled utilities:

$$L = u_{m,i}^* + .5(u_{m,1}^* - u_{m,i}^*)$$

If a_1 and a_i are the only two options, a_i will be preferred for values greater than L . Then the expected utility of a_i , given that it is preferred to a_1 , can be obtained simply by integrating from L and normalizing:

$$u^*(E_n, z^{(n)}, a_i) = \int_L^\infty x f_i(x) dx \quad / \quad \int_L^\infty f_i(x) dx.$$

where $f_i(\cdot)$ is the postulated normal probability function for $u^*(E_n, \tilde{z}, a_i)$ given $z^{(n)}$. These integrals can either be approximated by a computer, or assessed from tables of normal probabilities and the unit normal loss integral (Raiffa and Schlaifer, 1961).

This assessment heuristic requires only one assessment in addition to those required for the auxiliary tree (Figure 3-4): the

variance σ^2 of the distributions of expected utility. This assessment may be most naturally obtained as a credible interval--e.g., the interval within which the expected utility is likely to fall with 95% probability--from which the variance may be computed. The result of the proposed heuristic is merely a lower bound on $u^*(E_n, z^{(n)}, a_i)$, however, both because the correlation was assumed to be -1, and because a_1 and a_i were assumed to be the only options.

The relations between σ^2 and c has not as yet been spelled out. However, both σ^2 and c should be chosen so that:

$$P(a_i | E_n, z^{(n)}) \leq \int_L^\infty f_\pi(x) dx$$

since the integral on the right overestimates the probability that a_i will be preferred.

3.4 Non-Optimal Behavior

When information events are omitted from a traditional VOI model of a decision problem, the decision maker's choices may be inconsistent with that model. In another class of cases, however, the consistency condition (see Section 2.3.1) may fail not because the model is incomplete, but because the decision maker fails to act optimally on the information which he has. The reasons for such inconsistency with the model may include, on the one hand, motivation, fatigue, and constraints on time, memory, and attention, or, on the other hand, limitations of knowledge (see Section 2.3.2.1).

An important consideration in C^2 design, particularly in the provision of inference and decision aids, is the extent to which they will be used or--on account of habit, distraction, or distrust--ignored. Clearly, in the context of system design, the "consistency" of the user with the designer's model will reflect the compatibility of the system with the perceptual, cognitive, and organizational demands of the situation, i.e., it will reflect the designer's skill.

When, for whatever reason, the decision maker fails to maximize expected utility as modeled by the designer, the acts as events technique is appropriate. We can use the choice equation and an auxiliary decision tree, based on standard VOI analysis, to assess the probabilities for actions.

As before, the parameter c measures "consistency": the expected agreement of the decision maker with the system designer's model (depicted in the auxiliary tree). However, in this case, we are assessing not the completeness of the model, but the quality of the user's decision processes. When c equals zero, the model is unable to predict behavior, since choices are wholly unaffected by the considerations the designer regards as critical. (It may, however, be governed by response biases, represented by the b_i .) As c approaches infinity, behavior comes increasingly under the rational control of the modeled information events.

A plausible rationale for the application of the choice equation to acts as events is the assumption that the expected utility of an action, $u^*(e_n, z^{(n)}, a)$, as assigned by the system designer, stands in for a random variable. This variable represents the expected utility assigned by the decision maker to the act. In the case of incompletely modeled information events, uncertainty is due to the effect of such events, as yet unknown, on the decision maker's assessment. In the case of non-optimal behavior, however, uncertainty concerns the degree of closeness of the decision maker's expected utilities to what we assume are the "true" (i.e., the modeled) expected utilities.

Note that in the application of acts as events to non-optimal behavior, there is no need to adjust the expected utilities, $u^*(e, z, a)$, for actions. Unlike the case of omitted information events, no presumption exists that a given act has been chosen for a reason. Thus, the lower c falls, the more degraded the expected utility of the decision problem with information, $u^*(e)$.

3.4.1 Application to system design. As noted above, although c reflects the user's decision processes, it will, to a large extent, be dependent on properties of the C^2 system being evaluated. It is a measure of system usability. Thus, a goal of system design is--other things being equal--maximization of the value of c . Factors affecting c include such diverse considerations as legibility of displays, provision of inference aids to draw out the implications of raw data, and appropriate training in the use of the system.

In some cases, design questions may involve a tradeoff between improvements in usability and improvements in "pure informational" value. The latter is the difference between the modeled expected utilities:

$$EVSI_m(e_n) = u_m^*(e_n) - u_m^*(e_o)$$

i.e., the value of the decision problem with and without experiment e_n , given that the information is used optimally. (It may be computed by standard VOI analysis, as in the auxiliary tree of Figure 3-5.) A tradeoff might occur, for example, when the algorithm which is optimal for solving a given problem in a purely technical sense, is less likely to be used than a heuristic which is more familiar or faster to apply.

The methods outlined above could be used to locate the optimal point on such a tradeoff. Design options may be scored on the two dimensions, usability (c), and pure information value ($EVSI_m(e_n)$). A region of feasible options and an "efficient frontier" of undominated options may be defined within this two-dimensional space. The options which are not eliminated in this way can be evaluated by reference to the expected value of information, $EVSI(e_n)$. For each point in the space, this quantity, in accordance with the acts as events technique and the choice equation, uses the assessment of c to degrade the pure information value and arrive at a measure of overall utility. "Isopreference"

contours may be defined which represent the combinations of usability and of pure information value which are equivalent in terms of overall utility. Standard techniques can then be used to locate the best feasible design option.

3.4.2 Application to value of analysis. The evaluation of decision analysis (see Section 2.3.3) raises issues similar to those involved in non-optimal behavior, since there may be a discrepancy between the expected utilities assessed by a decision maker and the "true" beliefs and desires. The critical difference is that in the evaluation of decision analysis the true expected utilities (those that would arise in "perfect" analysis) are unknown. Thus, we compare the assessed expected utilities with the conditional expectation of the perfect expected utilities.

We construct an auxiliary tree, as in Figure 3-5, in which the information event is the outcome of imperfect analysis (i.e., the vector of expected utilities assigned to options); an uncertainty node after the decision node represents the possible outcomes of perfect analysis, conditional on the outcome of imperfect analysis. Since perfect analysis is assumed to reflect the decision maker's true beliefs and desires, optimal behavior consists in rolling back this decision tree, and selecting the act which maximizes the expected value of expected utility according to perfect analysis, conditional on imperfect analysis.

The auxiliary tree reflects, in essence, the proposal of Nickerson and Boyd (1978), that the output of decision analysis be treated as information within a standard VOI framework, and used to update expectations concerning "true" expected utilities. Although perfect analysis is unobtainable, this model describes how the decision maker may still behave optimally, with respect to his second-order assessments of the validity of imperfect analysis.

The behavior of the less than optimal decision maker can be predicted, by means of the choice equation, from the auxiliary tree together with an assessment of c . c measures the agreement of the expected utilities which the decision maker assigns to actions with the conditional expected values of perfect expected utilities. Thus, it reflects the extent to which the decision maker performs the required second-order assessments. When c approaches infinity, we have the Nickerson and Boyd model in which imperfect analysis is assessed in the light of prior beliefs concerning its relevance to perfect values. For c equal to zero, behavior is random with respect to these assessments.

By parameterizing consistency in this way, we avoid assuming either that decision makers are always consistent (Nickerson and Boyd) or, on the other hand, that they always act unquestioningly on the recommendation of the analysis, however bad (Watson and Brown, 1975).

3.5 An Inferential Structure for VOI

When acts are treated as events, they may be incorporated as variables within a probabilistic inference structure. Such a representation may possess advantages, in terms both of assessment and computation, over the traditional decision tree. One further condition must be satisfied, however, if these advantages are to be realized.

In Figure 3-2, we assumed that the system designer was evaluating a data base, E_n , and wished to model explicitly only selected experimental outcomes, $z^{(n)}$. Probabilities for the modeled states of the world, s , were dependent on the actual values of $z^{(n)}$. However, unmodeled events might be predictive for s as well. We now consider the special case in which s is independent of all unmodeled information events.


3.5.1 Inference tree representation. If this independence holds, we can substitute an alternative form of representation for Figure 3-2, in terms of hierarchical inference (HI). (See special issue of Organizational Behavior and Human Performance, December, 1973.) Hierarchical inference involves a target variable, input variables, and intermediate variables. The object is to obtain a probability distribution over values of the target variable, given particular values of the input variables. It may be easier to proceed by assessing probabilistic relationships between input and intermediate variables and between intermediate variables and the target variable; if so, HI enables us to derive the desired target distribution.

Figure 3-5 depicts the assessment problem in the form of an "inference tree" (Kelly and Barclay, 1973). Each node corresponds to a variable, and one node is shown below another if knowledge of the lower variable is required for an inference about the higher variable. The object of the inference in Figure 3-6 is to derive a probability distribution over terminal utility, u , for particular values of E (e.g., E_n, E_{n-1} , or e_0) and $z^{(n)}$. This distribution will be marginal with respect to intermediate variables: acts, a , states of the world, s , and the vector of expected utilities, u_m^* , from an auxiliary tree.

3.5.2 Computation of VOI. Once the distribution on u is obtained, we can proceed to derive the same quantities that figured in the more traditional Acts as Events evaluation procedure of Section 3.2 (See Figure 3-2). First, by taking the expected value of the distribution on u , we get the expected utility of the subsequent decision, after observing the results of the experiments in E , but marginal over all outcomes except $z^{(n)}$:

$$u^*(E, z^{(n)}) = E_{u|E, z^{(n)}} u.$$

Figure 3-6
Inference Tree



Next, we compute the expected utility of the decision problem with the information in E , by taking the expectation with respect to the modeled experimental outcomes, $z^{(n)}$:

$$u^*(E) = E_{z^{(n)}} | E u^*(E, z^{(n)}).$$

Once we have done this for all the relevant values of E , we can easily compute VOI; e.g.,

$$EVSI(E_n | E_{n-1}) = u^*(E_n) - u^*(E_{n-1}).$$

3.5.3 Assessments on u ; independence assumptions. What is distinctive about the inferential approach, of course, is not the foregoing calculations, but the derivation of the distribution on u from Figure 3-6. Two features of the inference tree are critical (cf., Kelly and Barclay, 1973): (i) it contains no closed paths. (The duplication of an input variable, $z^{(n)}$, presents no problems, however.) (ii) the probability of any variable, conditional on the variables immediately below it, is independent of any other variable in the tree.

Given these conditions, we decompose the assessment of probabilities on u as follows:

$$(1) P(u | E, z^{(n)}) = \int_s \int_a P(u | a, s) P(a | E, z^{(n)}) P(s | z^{(n)})$$

where \int is the general summation operator. Equation (1)

allows us to draw a sharp line between two roles of information: (1) as a conditioning event in the system designer's assessment of probabilities for states of the world, and (2) as a conditioning event for the system designer's assessment of probabilities for the decision maker's selection of actions. These roles are distinct because we do not wish to assume the consistency condition: that the decision maker selects actions

ly maximizing expected utilities based on probabilities assessed by the system designer. These probabilities are, however, still relevant to the evaluation - since states of the world affect the utility payoff for a given action.

Note that equation (1) depends on certain independence assumptions:

$$P(u|E, z^{(n)}, a, s) = P(u|a, s),$$

$$P(s|E, z^{(n)}, a) = P(s|z^{(n)}).$$

The first of these is implied by condition (ii) above, and is justified by the decomposition of total utility into the cost of the experiment and other aspects of utility (Section 2.2.4).

The second independence assumption has two parts. The Independence of s and a , conditional on E and $z^{(n)}$, reflects conditions (i) and (ii), and is by far the more important. Such independence holds if, as we assumed at the outset of this section, no unmodeled information events bear on the occurrence of s . In that case, the assessment of probabilities for s need not take into account the impact of unmodeled events on action.

The independence of s and E , conditional on $z^{(n)}$, is less important. This says that whether or not an experiment is included in a data base has no effect on the probabilities assessed by the system designer conditional on its (hypothetical) outcome. (We assume, with Raiffa and Schlaifer, 1961, that Z contains sufficient information to identify the relevant experiment in E .) This is based on the assumption that experiments do not affect states of the world (Section 2.2.6). It is not critical for the HI approach that this be true, however.

Both aspects of equation (1) - the effect of information on probabilities of actions and on probabilities of states of the

world - may be further decomposed. The assessment of act probabilities is decomposed as follows:

$$P(a|E, z^{(n)}) = \int_{\underline{u}_m^*} P(a|\underline{u}_m^*) P(\underline{u}_m^*|E, z^{(n)}).$$

This introduces a further independence assumption, in accordance with condition (ii):

$$P(a|E, z^{(n)}, \underline{u}_m^*) = P(a|\underline{u}_m^*).$$

We have already described (Section 3.3.2) how act probabilities may be assigned, by means of the choice equation, as a function of \underline{u}_m^* - i.e., the vector of expected utilities for actions computed from an auxiliary model. The required independence is thus insured. Turning now to $P(\underline{u}_m^*|E, z^{(n)})$, expected utilities for actions, \underline{u}_m^* , are computed within the auxiliary model by standard procedures of rolling back the tree (Section 3.3.1). This step is therefore deterministic:

$$P(\underline{u}_m^*|E, z^{(n)}) = \prod_i \delta(u_{m,i}^* - u_m^*(e_n, z^{(n)}, a_i)),$$

where $\delta(x) = \begin{cases} 1, & \text{if } x = 0 \\ 0, & \text{otherwise.} \end{cases}$

The other limb of Figure 3-6, representing the predictive impact of information for states of the world, can be decomposed as well. $P(s|z^{(n)})$ can be assessed indirectly, by means of Bayes' Theorem, in terms of $P(z^{(n)})$ and $P(z^{(n)}|s)$ (Section 2.2.4). Moreover, additional intervening variables may be inserted between s and $z^{(n)}$ if the system designer feels they will improve his assessments, regardless of whether or not such variables would be known to the decision maker. Indeed, a more general form of the inference diagram is possible, in which the experimental outcome variable, z , ranges over different outcomes in the righthand and lefthand branches. Thus, the system designer

can decompose his assessments on s in one way and his assessments on a in quite another. He need not suppose that the decision maker possesses the same information, or prefers the same decomposition, as he does.

3.6 Unspecified States of the World: Credibility

We have seen how experimental outcomes can be omitted from VOI analysis by modeling acts as events (Section 3.2). Further simplification in the use of VOI to evaluate information systems can be achieved by omitting states of the world. In this section we discuss the conditions under which an acts as events model can facilitate this step.

In traditional preposterior analysis, since states of the world follow the decision node, they may, of course, be integrated out. It is hardly possible to model explicitly all the events which could affect utility. To the degree that important factors are neglected, however, the price of increased simplicity is a more holistic, and probably less credible, estimate of the expected utilities for actions. The impact of information on the inferences of the decision maker and the system designer are no longer modeled in any detail. The resultant increase in uncertainty about $u^*(e, z, a)$ creates uncertainty in subsequent conclusions. The evaluation of the decision problem with an experiment, $u^*(e)$, is affected since this is the expectation, with respect to experimental outcomes, of the maximum expected utility for actions. Similarly, the computation of value of information ($EVSI(e)$) becomes less credible.

If one wishes to define "credibility" formally, we may think of it, following Tani (Section 2.3.2.2), as the "closeness" of an obtained assessment to the "authentic" value. If a probability distribution on authentic values is assessed conditionally on the obtained estimate, "closeness" is inversely related to the variance of that distribution. The effect of the credibility of a component assessment (e.g., $u^*(e, z, a)$) on the credibility of VOI ($EVSI(e)$)

can be computed by the method of decomposed error analysis (DEA) described in Brown (1968).)

States of the world may be integrated out when acts are modeled as events, too (Figure 3-7). Once again, we can expect some reduction in the credibility of estimates for expected utilities of actions. In some instances, however, the sensitivity of $u^*(E)$ to $u^*(E,a)$ will be less in an acts as events model than in a comparable preposterior model (see the Technical Note appended to this report). The reason is that a weighted average of a set of random variables tends to be more precise than the maximum.

The exact degree of sensitivity of $u^*(E)$ to $u^*(E,a)$ depends in a complex way on the nature of the acts as events model. In particular, it depends on whether expected utilities are used to predict actions and, if so, how strong the relationship is.

We can - if we wish - continue to predict the decision maker's behavior by reference to the expected utilities of acts. It was shown in Section 3.3.2 how an auxiliary tree can be used to compute the u_m^* for that purpose. In this case, the effect on credibility of omitting states of the world is somewhat larger, since the weights for expected utilities (i.e., the act probabilities) are themselves functions of u_m^* . In some instances, sensitivity will be much greater than in a comparable preposterior analysis.

A difference between acts as events and preposterior analysis, of course, is that in the former we are not compelled to predict behavior using expected utilities of acts. Instead of employing an auxiliary tree and the choice equation, we may model the impact of information on action independently - for example:




Figure 3-7

Integrating out experimental outcomes and states of the world.

- by observing actual behavioral patterns in natural or experimental settings;
- by noting policy guidelines or prescribed doctrines;
- by using descriptive psychological theories of behavior.

If this type of modeling is feasible, the loss of credibility due to omitting states of the world is reduced (see Technical Note).

These results are consistent with the following intuitions: for acts as events, it seems advisable to omit important states of the world only if the analytical resources saved are shifted to some independent method of modeling act probabilities. Conversely, if there is such an independent method, the value of modeling states of the world is reduced. We turn now to some techniques based on the supposition of independent modeling of act probabilities.

3.7 Utility Swing

A convenient formulation of many VOI problems is in terms of opportunity loss, or the cost of errors (Section 2.2.3, 2.2.7, and 2.2.10.1). This approach is particularly appealing when we have decided not to model states of the world, since direct assessment of utility differences is only one step beyond direct assessment of the utilities themselves. It turns out, however, that an application of the opportunity loss concept within an acts as events context is not quite as straightforward as it might appear.

The natural way to proceed would be as follows, with reference to the diagram of Figure 3-7:

$$EVSI(E_n) = u^*(E_n) - u^*(e_o)$$

$$= E_{a_E|E_n} u^*(E_n, a_E) - E_{a_o|e_o} u^*(e_o, a_o)$$

$$= E_{a_o|e_o} E_{a_E|E_n} [u^*(E_n, a_E) - u^*(e_o, a_o)]$$

where we have used a_E and a_o to refer to the informed and uninformed acts, respectively. With this approach we need to assess probabilities over the options both for the decision maker's initial preferences, a_o , and for his choices, a_E , after observing the data base contents, E_n . The quantity in brackets represents the utility swing for a particular combination of initial preference and informed choice, and may be assessed directly. (Note that we could also define opportunity loss in terms of perfect information and let:

$$EVSI(E_n) = l^*(e_o) - l^*(E_n)$$

$$= [u^*(PI) - u^*(e_o)] - [u^*(PI) - u^*(E_n)]$$

(see Section 2.2.10.1). We have chosen instead to regard "errors" as differences made in behavior by ignorance of a particular data base, E_n , and not by lack of "perfect knowledge". However, parallel considerations would apply to the alternative treatment.)

To see the problems that arise in the context of acts as events, note that with the standard opportunity loss treatment, when information does not affect choice, utility swing is zero (Section 2.2.7). The crucial intuition underlying our assessments is that we are evaluating the costs of errors, what we would have gained by switching responses. On account of our assumption that experiments do not affect states of the world (Section 2.2.6), we would therefore expect to find that:

$$u^*(E_n, a_E) = u^*(e_o, a_o)$$

when $a_E = a_o$.

This however, is not the case with acts as events. Utility assessments are conditional upon the actions performed, and upon the amount of knowledge one supposes to have informed the action (Sections 2.3.2.2, 3.3.4). Thus, even when the decision based on E_n is the same as would have been taken without it, we assess:

$$u^*(E_n, a_E) > u^*(e_o, a_o),$$

because the fact that the action is chosen on the basis of knowledge gives us more information about utility than the mere fact that it is chosen.

Fortunately, there is a formulation in terms of opportunity loss which fits the intuition that we should be concerned with errors. We can repartition the space of outcomes so as to establish the desired equality of expected utilities. All that is required is that we insert a chance node after the uninformed act, a_o , which corresponds, hypothetically, to what the informed act, a_E , would have been (Figure 3-8).

The expected utility assessment, $u^*(e_o, a_o, a_E)$, is now conditioned on the hypothesis that, in this situation, if the decision maker had known E_n , he would have chosen a_E . Thus, if the uninformed decision maker makes the same choice in that situation, he gets the same utility. The effect of this partitioning is that:

$$u^*(e_o, a_o, a_E) = u^*(E_n, a_E) \geq u^*(e_o, a_o)$$




Figure 3-8

when $a_E = a_O$, and

$$u^*(e_O, a_O, a_E) \leq u^*(e_O, a_O)$$

when $a_E \neq a_O$; i.e., expected utility is adjusted downward when we know that the informed choice would have been different.

We can now reformulate the acts as events model in terms of opportunity loss:

$$\begin{aligned} \text{EVSI}(E_n) &= u^*(E_n) - u^*(e_O) \\ &= E_{a_E|E_n} u^*(E_n, a_E) - E_{a_O|e_O} E_{a_E|e_O, a_O} u^*(e_O, a_O, a_E). \end{aligned}$$

And, assuming that the informed decision maker knows what his initial choice would have been, i.e., $P(a_E|e_O, a_O) = P(a_E) = P(a_E|E_n)$, we have:

$$\text{EVSI}(E_n) = E_{a_O|e_O} E_{a_E|E_n} [u^*(E_n, a_E) - u^*(e_O, a_O, a_E)].$$

The quantity in brackets can be interpreted legitimately in terms of opportunity loss. In effect, it compares a_O and a_E in the same situation. It corresponds to what the decision maker imagines he would pay to be allowed to change his mind if, after choosing a_O , he observes E_n , and as a result now prefers a_E . He will, of course, pay nothing to retract his decision if E_n does not cause his preference to shift.

The opportunity loss formulation is convenient, as noted in Section 2.2.10.1, when there is an additive component of utility common to all options a . Moreover, the required assessments seem quite natural. Still, it should be noted that in the acts as events context, the number of required utility swing assessments is actually larger (the square of number of options) than

the number of utility assessments that would ordinarily be needed (twice the number of options). Even when the uninformed choice, a_0 , is known with certainty, assessments are required for each possible switch from a_0 , i.e., for each option.

3.8 Unspecified Options

Thus far, every VOI technique which we have reviewed or proposed has required that the options confronting the decision maker be specified in advance. This is a quite unreasonable demand when information systems are expected to perform in changing and largely unpredictable environments. It is still less reasonable when some of the information to be evaluated provides assistance in identifying options. Finally, we can expect enormous economy of effort by framing a VOI analysis with unspecified options.

A natural starting point is the opportunity loss idea in the context of acts as events (Section 3.7). We now assume that the utility swing,

$$U(E_n) = u^*(E_n, a_E) - u^*(e_0, a_0, a_E),$$

is (roughly) constant for all $a_E \neq a_0$. Then we have:

$$\begin{aligned} EVSI(E_n) &\approx E_{a_0|e_0} E_{a_E|E_n, a_E \neq a_0} U(E_n) \\ &= U(E_n) \sum_{a_0} \sum_{a_E \neq a_0} P(a_0|e_0) P(a_E|E_n) \\ &= U(E_n) \sum_{a_0} \sum_{a_E \neq a_0} P(a_0 \& a_E | e_0), \end{aligned}$$

again assuming that $P(a_E|E_n) = P(a_E|a_0, e_0)$. Thus,

$$EVSI(E_n) \approx U(E_n) \cdot S(E_n)$$

where $S(E_n)$ is the probability that the decision maker would switch options, if he had knowledge of E_n , and $U(E_n)$ is the expected shift in utility given that he does switch options as a result of E_n .

Both $U(E_n)$ and $S(E_n)$ can be directly assessed, with a fair degree of naturalness, in a variety of circumstances. When this is possible, it is unnecessary to spell out the options between which the decision maker might switch. It is only required to assess the likelihood that the information will cause some change or other, and the expected benefit.

3.9 Decomposing Shift Probabilities: Information Value and Usability

For many purposes of C^2 system assessment it is desirable to distinguish features relating to "pure informational value" from those relating to "usability". Unfortunately, efforts to categorize evaluative criteria in such terms have seldom been rigorous or formally justified (see Section 2.1.2). We have suggested one approach to this problem based on an auxiliary decision tree and the choice equation (Section 3.4.1). We now outline another, much simpler method which flows from the opportunity loss conception.

The essential idea is to decompose the probability that the decision maker will change options due to E_n into two preconditions:

- the information content of E_n would suggest a change in options to an "optimal" decision maker;
- an actual user would change to the indicated option given that an optimal user would.

The suggested decomposition is summarized in Figure 3-9. The information system would cause an optimal decision maker to change from his initial preference, a_0 , to a different choice, a_1 , with probability K . Given that an optimal decision maker changes, the actual user will switch his preference, a_E , to the indicated option (or to one nearly as good) with probability K . If he does




Figure 3-9

Probability tree for information value and usability.

switch, his swing in expected utility is U , as defined above. If, on the other hand, he remains at his initial preference, or switches without apparent reference to the information system content, his expected utility swing is zero.

If the information system content suggests to an optimal user that he stay with his initial choice, a_0 , we assume - for simplicity - that the actual user will remain at a_0 with probability one. It seems unlikely that the decision maker would very often be deterred from his chosen path by the illegibility of a confirming information display, or even by a fallacious interpretation of it.

We thus have a "multiattribute" decomposition of $S(E_n)$ in terms of two parameters: information value, $R(E_n)$, and usability, $K(E_n)$:

$$S(E_n) = R(E_n) \cdot K(E_n).$$

This MAU model is not ad hoc (see Section 2.1.2). The multiplicative combination rule is a direct implication of our probabilistic assumptions; and the interpretation of R and K as probabilities provides a framework in which their meaning can be clarified and communicated.

An important part of such clarification is the definition of "optimal decision maker" (ODM). First, it should be clear that ODM is not omniscient. We attribute to him roughly the same substantive knowledge to be expected in the actual user of the information system. Otherwise, of course, the information provided by the system could never cause him to change his mind.

On the other hand, we assume that ODM is consistent (Section 2.3.1). This means that he rationally incorporates new information into his belief structure: he assesses probabilities for states of the world conditional on the information, assesses utilities for each

combination of act and state of the world, and selects an action by maximizing expected utility. A second assumption about ODM is that, in doing so, he is undeterred by such factors as fatigue, workload, inattention, failures of memory, or illegible displays.

R is the probability that ODM, so equipped, will switch options as a result of the provided information. $(1-K)$ measures the likelihood that inconsistency due to performance factors (fatigue, etc.) or errors of knowledge (e.g., the use of an incorrect decision rule) will negate this potential switch in actual practice.

3.10 Application to VOA

When inference or decision aids are provided by an information system, they may improve the "usability" of other information items. They may reduce the effects of workload and fatigue by automating certain functions, and they may correct errors of knowledge by helping to draw the implications of information for belief and action.

Decision aids, however, may possess information value in their own right, considered in isolation from other elements of the data base. Note that although ODM is assumed to be a consistent probability assessor, he is not necessarily an authentic one (Tani 1978; Section 2.3.3.2). That is, his assessments need not fully incorporate the relevant knowledge which he already possesses - and which, if he had "infinite time" for reflection, he would bring to bear on the assessments. A decision aid can cause an optimal decision maker, so defined, to change his mind. It can stimulate the generation of available options; and it can improve his assessment of probabilities by decomposing them in more natural ways, that draw on more of his knowledge. His judgments are consistent conditional on any given knowledge set.

Since ODM is consistent, he will evaluate the output of a decision

aid as "information", assessing its validity in the context of other beliefs, and acting accordingly (Nickerson and Boyd, 1980; Section 2.3.3.3). Note that a_i - the option indicated as optimal by this process - need not be the option explicitly recommended by the decision aid.

Actual decision makers, of course, may either fail to appreciate a good decision aid or act unquestioningly in accordance with a bad one (cf., Watson and Brown, 1975; Section 2.3.3.1). K, therefore, measures the "usability" of the decision aid - the extent to which its impact on behavior reflects a rational evaluative process on the part of the user.

The proposed technique satisfies two desiderata for the evaluation of decision analysis, discussed in Section 2.3.2:

- It does not assume that an actual decision maker is perfectly consistent in the probability judgments used to evaluate the analysis.
- It does not require the prior specification of options available to the decision maker.

Thus, in addition to being simpler than current approaches, it is applicable to decision aids which assist in the modeling of options.

4.0 INFORMATION EVALUATION FOR SYSTEM USERS

4.1 "Intelligent" Computer Aids for Information Selection

In this chapter we turn to the second level of information selection described in the Introduction: the "user dialogue". Certainly one of the most important characteristics of future C³I systems will be the ability to perform proactively and intelligently in the on-line provision of information to decision makers. Users should not have to traverse complex sequences of computer-generated menus or cope with complex thesauri of computer-recognizable terms in order to tell the system what information is needed. To whatever extent possible, the system itself should quickly determine the information needs of the particular user, and select optimal methods of data presentation based upon its knowledge of human factors and its internal determination of the expected value of information.

While system designers must consider the broad range of scenarios in which a system is to function, our concern shifts now to a particular user in a particular situation. Even though, in a sense, the user has less to consider, the constraints on the complexity of the user dialogue are no less severe than for the designer dialogue. The user of a C² system will be operating under heavy pressures of time and cognitive load, particularly in situations of combat. A determination of the information that is of greatest value must take place as quickly and effortlessly as possible. We may add that elaborate and complex programs to perform a full VOI analysis would not appear to be practicable adjuncts of data-base systems, given the present state of the art in computer aids.

In our discussion of the user dialogue, we will draw upon concepts developed in Chapter Three, particularly Sections 3.7 and 3.8, for a method which may prove both practicable and effective.

The method is based on the fundamental concept that information has a value only insofar as it causes a decision to be changed. In essence, then, we require the computer to ask only how likely data in a given category are to cause a switch in the preferred options, together with an estimate of how valuable such a switch would be. The data that are then presented to the user will be those which maximize the product of the probability of a switch and the expectation of the shift in utility. However, a direct assessment of these entities may not be possible for the decision maker (DM) without some more solid background against which to base these assessments. Thus, the procedure we advocate provides for a certain initial structuring of the decision problem. There are, however, many different levels of effort and complexity at which the assessments could be carried out. We shall indicate the advantages and disadvantages of the various levels of decomposition.

4.2 Outline of Proposed Program

The basic idea of the approach is shown in diagramatic form in Figure 4-1. In Step 1, the computer asks the DM to make a list of all those options that the DM considers to be potentially selectable. The purpose of this is to make the questioning in later steps less of an abstract exercise. The DM is then asked to select that option which he considers to be the best, in the absence of any information that might be provided by the data base.

In Step 2, the DM is asked by the computer to think about which areas of major uncertainty the DM would like most to have resolved in order to feel more happy about the option selected. It will be emphasized to the DM that consideration should be centered upon those uncertainties where potentially available information might cause the preferred option to switch. Thus, even were the DM

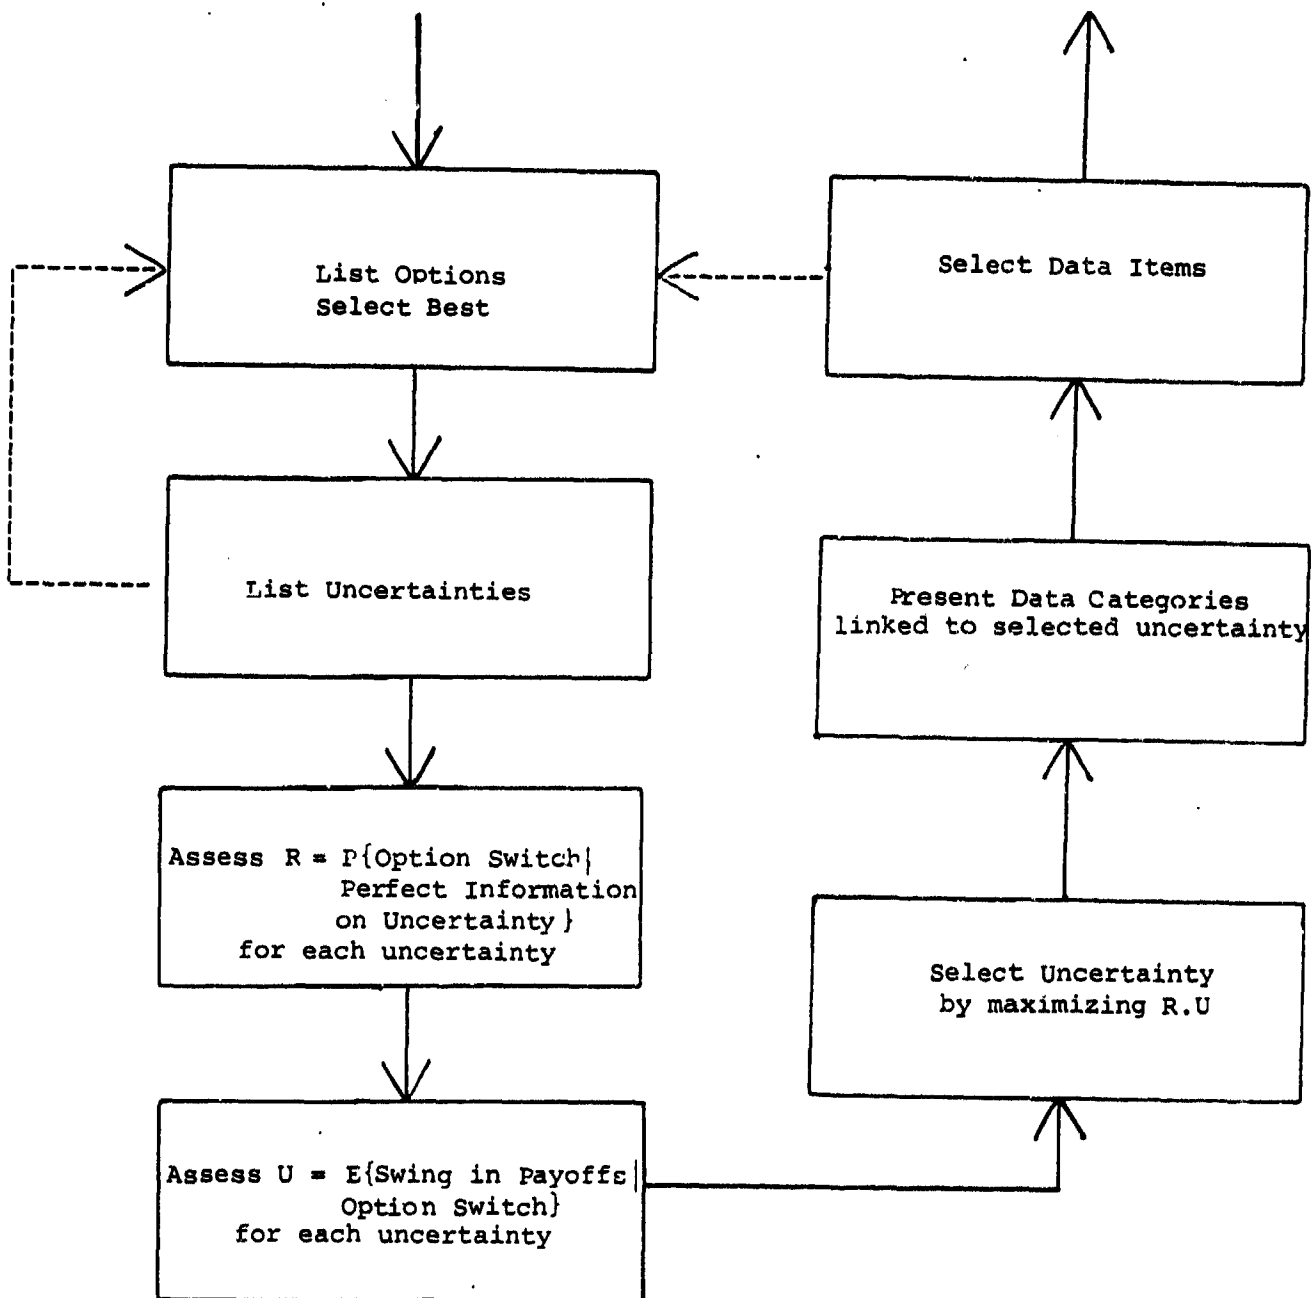


Figure 4-1

particularly worried about a certain area of uncertainty, if he did not feel that any potential resolution of that uncertainty would cause the preferred option to be switched, then the information would have (at that point and time) no value. Each individual type of uncertainty may, or may not, be explicitly linked via the computer to the relevant decision options. It is quite conceivable that, after producing such a list of uncertainties, the DM will want to return to Step 1 in order to include one or more previously unconsidered options. Such an iterative procedure should, of course, be built into the decision aid.

In Step 3, the computer will need to elicit from the DM the probabilities that the preferred option would be switched given perfect information on each of the uncertainties listed in Step 2 above. In Step 4, the computer would need to elicit the expected change in the value of the decision, given a switch due to perfect information, for each of the uncertainties. These changes in utility would perhaps be best assessed on a 0 to 100 scale, where 0 indicates the worst possible outcome and 100 represents the best possible outcome. After Steps 3 and 4, it will be possible to decide which uncertainty it is of most value to resolve, by multiplying together the probability of a switch given that resolution, and the expected change in utility from a switch. The most valuable is obtained by finding the uncertainty which maximizes this product.

4.3 Levels of Analysis

The steps described in the last paragraph are the heart of the process. However, many different levels of effort could potentially be programmed into the decision aid. We now indicate, in increasing order of effort, some of the potentially available levels of decision aid:

- The computer could simply ask directly of the DM which uncertainty would be of the most value to have resolved. Such a direct assessment would not, we suspect, provide much useful aid to the DM.
- One could assume that the expected utility from a switch in the preferred option was approximately the same over all uncertainties and simply ask the DM which uncertainty would, upon resolution, provide the greatest probability of a switch in options. That uncertainty would then be targeted as the one concerning which data should be presented. It will be noted that this approach does not require the assessment of numerical values for the probability of a switch--merely the elicitation of a maximal element. The degree of approximation inherent in the initial assumption of equi-value switches will, of course, depend on the given context. In certain situations, this approach may indeed be the best because of its simplicity.
- The analysis could be performed as indicated in the previous section, by assessing both the probability of a switch and the expected utility of a switch for each uncertainty. The assessment of the probabilities of switching could be done at a more or less refined level and similarly the expected switches in utility, although we assume that simple, direct assessments on a 0 to 100 scale (perhaps via ratio judgments) are the most practicable.
- One could build a full multiattribute utility (MAU) model instead of Step 4. The DM would then be required to indicate how much gain (or loss) on each of various key criteria could be expected from the potential switches in preferred options.

- The assessments could be conditional upon which particular switch in preferred option was made. Thus, the computer would assess from the DM the probability of switching to each of the options listed in Step 1 and then assess the expected switch in utility individually for each possible option.
- The greatest degree of complexity would be to build a full decision-analytic model of the problem and to perform a complete VOI calculation thereon. However, as indicated at the beginning of this section, such a procedure is not at present to be contemplated.

4.4 Presentation of Data

At this stage, the computer will have gained from the DM knowledge of the area of uncertainty to which presented data should refer. It is now necessary to decide which individual items of data in fact pertain to this uncertainty and also to determine a prioritization of all such items of data. It appears to be necessary that some pre-sorting of the items in the data base should have been carried out prior to this analysis. At the simplest level, one could simply associate with each item of data several key words. A list of key words could then be presented to the DM, on whom would fall the chore of linking the selected area of uncertainty with presented key words. This, of course, would be the most basic form of information selection, and one upon which we believe it to be fairly easy to improve, given a certain amount of ingenuity and pre-modeling effort.

The key to such more sophisticated methods of choosing data is that, for a given data base, the types of uncertainty that may be faced by decision makers will be finite and of small number.

Pre-modeling efforts are required to have associated with each item of data in the data base an indication of which type of uncertainty is addressed by that item. This indication could be either in a binary yes or no form, or, in a more sophisticated version, by a number between, say, 0 and 10, indicating the degree of relevance of that information to each type of uncertainty. Once the important uncertainty had been isolated by Steps 1 through 5, any of a variety of algorithms could be used in order to decide which items of data should be presented to DM. As an example, we have the following:

- If the binary categorization has been used, the computer could simply present all those items which were classified as being relevant to the given uncertainty.
- If a more complete numerical scale was used, the computer could present items in order of their relevance to that given area of uncertainty.
- The statistical technique of cluster analysis could be used to find items of data that were similar in characteristics to a given prototype piece of data. This prototype could be an item of data impinging purely upon the selected area of uncertainty, though it could also be a more sophisticated version given to the computer by the DM, i.e., the computer might ask the DM, after having displayed the area of uncertainty of greatest importance, to feed in a "relevance profile" which the DM considered to be the most appropriate type of data that could be presented at that time.
- A more sophisticated form of cluster analysis could be used, based either on stochastic clustering or fuzzy clustering (see Appendix) which would permit the computer to "know" to what degree the given data element belonged to the relevant cluster. This would again provide for a prioritization of the presentation to the DM.

After presentation to the user of each item of data, the computer should present to the DM the possibility of going back to stage one. In this way, the DM would be able to add to his previous list of options any new ones that had occurred to him after

presentation of the data and also any other major uncertainties which now were worrying him. A new preferred option could be selected as required, new probabilities and utilities could be assessed, and, thus, a new primary area of uncertainty could be distinguished. The procedure would thus be iterative, as indicated in Figure 4-1 until the decision maker either felt sufficiently confident of the preferred option, or until available time and resources had run out.

5.0 TECHNICAL NOTE: CREDIBILITY

In this note we show how the credibility of the evaluation of a decision problem with information, $u^*(E)$, depends on the credibility of the assessments for the expected utilities of actions (Section 3.6). We will compare two versions of acts as events in this respect: one, using expected utilities to estimate probabilities for actions; the other, estimating those probabilities independently. Preposterior analysis will be used as a benchmark. The ultimate purpose is to assess the consequences of omitting states of the world in these types of models.

We use the choice equation to derive:

$$\begin{aligned} u^*(E) &= \sum_a P(a|E) u^*(E, a) \\ &= \sum_i \frac{u_{m,i}^* c}{\sum_j u_{m,j}^* c} (u_{m,i}^* + A_i) \end{aligned}$$

It is necessary to note several simplifications. First, we treat $u^*(E, a)$ as $u_{m,i}^*$ plus an adjustment A . A is zero for preposterior analysis ($c \rightarrow \infty$). For acts as events it is always positive, and may vary with u_m^* and c . In this analysis we will ignore A - in effect assuming that it is not very sensitive to the value of u_m^* . Nonetheless, for this reason we may underestimate the dependence of the variance $V(u^*(E))$ on $V(u_m^*)$ for acts as events.

Second, we are ignoring modeled information events, since the effect of uncertainty in the assessments of $P(z|E)$ is not our present concern, and such an effect would be the same for acts as events and for preposterior analysis when they model the same outcomes.

Finally, we assume that our uncertainty concerning the "authentic" values of the $u_{m,i}^*$, given the assessed ones, is constant across acts.

On these assumptions, using the method of Brown (1968), and writing P_k for $\frac{u_{m,k}^{*c}}{\sum_j u_{m,j}^{*c}}$, we get:

$$V(u^*(E)) = V(u_m^*) \sum_i \left(\frac{\partial \left[\sum_k P_k u_{m,k}^* \right]}{\partial u_{m,i}^*} \right)^2 + \text{H.O.T.}$$

$$= V(u_m^*) \left\{ \sum_i \left[c P_i (1 - P_i) + P_i - \frac{c u_{m,i}^{*c-1} \sum_{j \neq i} u_{m,j}^{*c+1}}{\left(\sum_j u_{m,j}^{*c} \right)^2} \right]^2 \right\} + \text{H.O.T.}$$

By ignoring the higher interactions of the u_m^* we do not affect the dependence of $V(u^*(E))$ on $V(u_m^*)$.

According to this equation,

$$V(u^*(E)) = K(c) V(u_m^*).$$

For traditional preposterior analysis our uncertainty concerning the decision problem with E is the same as our uncertainty about the alternative for which expected utility is maximum:

$$V(u^*(E)) = V(u_m^*).$$

For acts as events, $K(c)$ may be greater or less than 1 depending on c , the number of options, and the values of the u_m^* . Thus, when c equals zero, the rate at which $V(u^*(E))$ increases due to increases in $V(u_m^*)$ is less, in comparison to preposterior analysis:

$$V(u^*(E)) = V(u_m^*) \sum_i P_i^2 = V(u_m^*) / q$$

where q is the number of options (and $P_i = 1/q$ since $c = 0$). Here, $K(0) = 1/q$.

On the other hand, for $c = 1$,

$$\begin{aligned} V(u^*(E)) &= V(u_m^*) \cdot \sum_i [2P_i - \sum_i P_i^2]^2 \\ &= V(u_m^*) \cdot q \cdot (\sum_i P_i^2)^2. \end{aligned}$$

$\sum_i P_i^2$ has its minimum value ($1/q$) when each P_i equals $1/q$, i.e., when the $u_{m,i}^*$ are equal; it has its maximum value (1) when only a single response has expected utility not equal to zero. Thus,

$$q \geq q(\sum_i P_i^2)^2 \geq 1/q.$$

Hence, $K(1)$ may be considerably greater than 1 when options have very disparate expected utilities.

We turn now to the independent modeling of act probabilities, in an acts as events model. In this case,

$$V(u^*(E)) = V(u_m^*) \sum_i P_i^2 + \sum_i V(p_i) u_{m,i}^{*2} + \text{H.O.T.}$$

where $P_i = P(a_i|E)$. Again, $\sum_i P_i^2$ is minimum for equiprobable acts and maximum for exclusive preference. Thus,

$$1 \geq \sum_i P_i^2 \geq 1/q.$$

With independent modeling of act probabilities, therefore, the dependence of $V(u^*(E))$ on $V(u_m^*)$ is reduced in comparison to acts as events with $c = 1$ - and, subject to the qualification concerning the adjustment A , may be considerably less than in the preposterior analysis.

6.0 REFERENCES

Alberts, D. S. C² assessment: A proposed methodology (MTR-80W00041). McLean, VA: The MITRE Corporation, January 1980.

Andriole, S. J. Another side to C³: Computer-based information, decision, and forecasting systems. Paper prepared for the Armed Forces Communications and Electronics Assoc. (AFCEA). Western Conference/Exposition, Anaheim, California, January 9-11, 1980.

Barclay, S., Brown, R. V., Kelly, C. W., Peterson, C. R., Phillips, L. D., and Selvidge, J. Handbook for decision analysis. McLean, VA: Decisions and Designs, Inc., September 1977.

Barclay, S., Chinnis, J. O., and Minckler, R. D. A prototype Projectile Tracking Radar (PTR) evaluation model, Part I (Technical Report 75-14). McLean, VA: Decisions and Designs, Inc., November 1975. (NTIS No. AD B010604L)

Barnes, M. J. Review of five military decision aids. China Lake, CA: Naval Weapons Center, May 1980.

Beard, L. H. Planning a management information system: Some caveats and contemplations. Financial Executive, May 1977, pp. 34-39.

Brown, R. V. Research and the Credibility of Estimates. Homewood, IL: Richard D. Irwin, Inc., 1968.

Brown, R. V. Modeling subsequent acts for decision analysis (Technical Report 75-1). McLean, VA: Decisions and Designs, Inc., May 1975. (NTIS No. AD A018888)

Brown, R. V. Heresy in decision analysis: Modeling subsequent acts without rollback. Decision Sciences, 1978, 9, 543-554.

Brown, R. V., Kahr, A. S., and Peterson, C. R. Decision analysis for the manager. New York: Holt, Rinehart, and Winston, 1974.

Brown, R. V., and Watson, S. R. Case studies in the value of decision analysis (Technical Report 75-10). McLean, VA: Decisions and Designs, Inc., October 1975. (NTIS No. AD A019263)

Chinnis, J. O., Kelly, C. W., Minckler, R. D., and O'Connor, M. F. Single Channel Ground and Airborne Radio System (SINCGARS) evaluation model (Technical Report DT/TR 75-2). McLean, VA: Decisions and Designs, Inc., September 1975. (NTIS No. AD A033094)

Cleverdon, C. W. Report on the testing and analysis of an investigation into the comparative efficiency of indexing systems (ASLIB Cranfield Research Project). Cranfield, England: College of Aeronautics, 1962.

Cohen, M. S., and Brown, R. V. Decision support for attack submarine commanders (Technical Report 80-11). Falls Church, VA: Decision Science Consortium, Inc., 1980.

Cooper, W. S. The paradoxical role of unexamined documents in the evaluation of retrieval effectiveness. Information Processing and Management, 1976, 12, 367-375.

Daly, J. A., and Andriole, S. J. Problems of applied monitoring and warning: Illustrations from the Middle East. Jerusalem Journal of International Relations, in press.

Davis, G. B. Management information systems. New York: McGraw Hill, 1974.

Decision Science Consortium, Inc. Conceptual decision-analytic models for use of fire related information in forest planning and escaped fire situation analysis. Falls Church, VA: Decision Science Consortium, Inc., December 1980.

Decisions and Designs, Inc. Early warning and monitoring prototype system: Sample output. McLean, VA: Decisions and Designs, Inc., January 1978.

Freeling, A. N. S. Decision analysis and fuzzy sets. Unpublished masters dissertation, Cambridge University, 1979.

Freeling, A. N. S. Fuzzy sets and decision analysis. IEEE Transactions on Systems, Man, and Cybernetics, 1980, SMC-10, 341-354.

Freeling, A. N. S., and Seaver, D. A. Decision analysis in Forest Service planning: Treatment of Mountain Pine Beetle (Technical Report 80-8). Falls Church, VA: Decision Science Consortium, Inc., September 1980.

Gallagher, C. A. Perceptions of the value of management information system. Academy of Management Journal, 1974, 17(1), 46-55.

Gorry, G. A., and Scott Morton, M. S. A framework for management information systems. Sloan Management Review, Fall 1971, pp. 55-70.

Herner, S., and Snapper, K. J. The application of multi-criteria utility theory to the evaluation of information systems. Journal of the American Society for Information Science, 1978 29, 289-296.

Huber, G. P. Organizational information processing systems: Logistical determinants of performance and behavior. Madison: University of Wisconsin, Graduate School of Business, May 1980.

Huber, G. P. Organizational information processing: Changes in the form and meaning of messages. Madison: University of Wisconsin, Graduate School of Business, July 1980.

International Public Policy Research Corporation. Early warning and monitoring system: Sample output. McLean, VA: International Public Policy Research Corporation, June 1979.

Kahneman, D., and Tversky, A. Intuitive prediction: Biases and corrective procedures (Technical Report PTR-1042-77-6). Eugene, OR: Decision Research, June 1977.

Keeney, R. L., and Raiffa, H. Decisions with multiple objectives: Preferences and value tradeoffs. New York: Wiley, 1976.

Larcker, D. F., and Lessig, V. P. Perceived usefulness of information: A psychometric examination. Decision Sciences, 1980, 11, 121-133.

LaValle, I. H. On cash equivalents and information evaluation in decisions under uncertainty. American Statistical Association Journal, 1968, 63, 252-290.

LaValle, I. H. On value and strategic role of information in semi-normalized decisions. Operations Research, 1980, 28, 129-138.

Lindley, D. V., Tversky, A., and Brown, R. V. On the reconciliation of probability assessments. Journal of the Royal Statistical Society, Series A, 1979, 142(2), 146-180.

Luce, R. D. The choice axiom after twenty years. Cambridge: Harvard University, Department of Psychology and Social Relations, June 1977.

Luce, R. D. Individual choice behavior. New York, Wiley, 1959.

Luce, R. D. A threshold theory for simple detection experiments. Psychology Review, 1963.

Mahoney, R. B., and Clayberg, R. P. Analysis of the Soviet crisis management experience: Technical report. Arlington, VA: CACI, Inc., September 1978.

Miller, R. S. A framework for establishing the need for and value of C² programs. Monterey, CA: U.S. Naval Postgraduate School, June 1980.

Mintzberg, H. Impediments to the use of management information. New York: National Association of Accountants, ____.

Nickerson, R. C., and Boyd, D. W. The use and value of models in decision analysis. Operations Research, 1980, 28, 139-155.

O'Connor, M. F., and Edwards, W. On using scenarios in the evaluation of complex alternatives. In Snapper, K. J. (Ed.), Studies in the analysis of complex decisions. Washington, D.C.: Information Resources Press, in press.

Raiffa, H. Decision analysis. Reading, MA: Addison-Wesley, 1970.

Raiffa, H., and Schlaifer, R. Applied statistical decision theory. Cambridge: The M.I.T. Press, 1961.

Samet, M. G., and Davis, K. B. Computer-based supervisory system for managing information flow in C³ systems: Pacing model. Woodland Hills, CA: Perceptronics, Inc., March 1977.

Samet, M. G., Weltman, G., and Davis, K. B. Application of adaptive models to information selection in C³ systems. Woodland Hills, CA: Perceptronics, Inc., December 1976.

Shannon, C. E., and Weaver, W. The mathematical theory of communication. Urbana: University of Illinois Press, 1949.

Shlaim, A. Failures in national intelligence estimates: The case of the Yom Kippur War. World Politics, 1976, XXVIII, 348-80.

Slovic, P., and Lichtenstein, S. Comparison of Bayesian and regression approaches to the study of information processing in judgment. Organizational Behavior and Human Performance, 1971, 6, 649-744.

Smith, W. A. Effectiveness of information systems (Report IE-14-7202). Bethlehem, PA: Lehigh University, Department of Industrial Engineering, June 1972.

Spector, B. I., Mahoney, R. B., Fain, J., and McIlroy, J. J. A crisis problem analyzer for crisis management: Technical report. Arlington, VA: CACI, Inc., September 1978.

Steeb, R., Chen, K., and Freedy, A. Adaptive estimation of information values in continuous decision making and control of remotely piloted vehicles. Woodland Hills, CA: Perceptronics, Inc., August 1977.

Swanson, E. B. Management and information systems: Appreciation and involvement. Management Science, 1974, 21, 178-188.

Tani, S. N. A perspective on modeling in decision analysis. Management Science, 1978, 24, 1500-1506.

Tani, S. N. Modeling and decision analysis (EES-DA-75-3). Stanford, CA: Stanford University, Department of Engineering-Economic Systems, June 1975.

Von Trees, H. L. Keynote address - Conference/Workshop on Quantitative Assessment of the Utility of Command and Control Systems, January 1980.

Watson, S. R., and Brown, R. V. The valuation of decision analysis. Journal of the Royal Statistical Society, Series A, 1978, 141, 69-78.

Watson, S. R., and Brown, R. V. Issues in the value of decision analysis (Technical Report 75-9). McLean, VA: Decisions and Designs, Inc., November 1975. (NTIS No. AD A019339)

Watson, S. R., Brown, R. V., and Lindley, D. V. Three papers on the valuation of decision analysis (Technical Report 77-2). McLean, VA: Decisions and Designs, Inc., May 1977.

Watson, S. R., Weiss, J. J., and Donnell, M. L. Fuzzy decision analysis. IEEE Transactions on Systems, Man, and Cybernetics, 1979, SMC-9, 1-9.

Yovits, M., Rose, L., and Abilock, J. Development of a theory of information flow and analysis. In J. Tow (Ed.), Advances in Information Systems Science. New York: Plenum, 1969-78.

Zadeh, L. A. Fuzzy sets as a basis for a theory of possibility. Fuzzy Sets and Systems, 1978, 1, 3-28.

Zani, W. M. Blueprint for MIS. Harvard Business Review, 1979, 48(6), 95-100.

Zmud, R. W. An emperical investigation of the dimensionality of the concept of information. Decision Sciences, 1978, 9, 187-95.

APPENDIX

FUZZY SETS AND INFORMATION SYSTEMS

INTRODUCTION

Fuzzy set theory is a relatively new discipline. The seminal paper was written by Zadeh in 1965, and since then a great deal of time and effort has been spent developing the concepts. In this section we shall present an overview of the foundations of the subject, and show how their application could be of potential value in the design and use of information systems.

An Overview of the Fundamentals of the Fuzzy Set Theory

The basic aim of fuzzy set theory is to handle imprecision. To overcome the need for precision, Zadeh (1965) argued the need for a new, fuzzy approach to the analysis of systems, and to this end he introduced his fuzzy set theory and the related concept of fuzzy logic. These ideas are compelling enough to have stimulated considerable study over the last thirteen years. In this section we present only a very quick treatment of the theory. For further detail see Watson, Weiss and Donnell (1979), Freeling (1979, 1980) or any of Zadeh's own papers.

Our treatment follows that of Watson et al. The central concept of fuzzy set theory is the membership function, which represents numerically the degree to which an element belongs to a set. This function takes on values between 0 and 1, and it is an extension of the idea of a characteristic function for a set. The membership function is assessed subjectively in any instance, small values representing a low degree of membership, and high values representing a high degree of membership. Freeling (1980a) discusses further the question of elicitation.

The calculus of fuzzy sets is based on three reasonable propositions which numbers of this type ought to satisfy:

- (a) The degree to which a belongs to both A and B is equal to the smaller of the individual degrees of membership.

$$\text{i.e., } \mu_{A \cap B}(a) = \min(\mu_A(a), \mu_B(a))$$

(using an obvious notation)

- (b) The degree to which a belongs either to A or to B is equal to the larger of the individual degrees of membership.

$$\text{i.e., } \mu_{A \cup B}(a) = \max(\mu_A(a), \mu_B(a))$$

- (c) The degree to which a belongs to (not A) is one minus the degree to which a belongs to A .

$$\text{i.e., } \mu_{\bar{A}}(a) = 1 - \mu_A(a).$$

Many possible uses for these concepts and this calculus in systems analysis have been suggested. One particular idea is to extend the notion of a function to allow fuzzy inputs and produce fuzzy outputs. To deduce the "fuzz" on the output given the "fuzz" on the inputs, we use the relationships above in the following way:

Let $\mu_i(x_i)$ be the degree to which x_i belongs to the possible set of numbers for the i^{th} input variable, and $\mu_0(y)$ be the degree to which y belongs to the set of possible numbers for the output variable. Then application of the rules above gives

$$\mu_0(y) = \text{Max}_{y=f(x)} [\text{Min}(\mu_1(x_1), \mu_2(x_2), \dots, \mu_n(x_n))]$$

(1)

The idea of Equation (1) is that for each "possible" set of values of the inputs (i.e., for all sets of values having non-zero membership functions) the output value is calculated and then a membership grade is given to this output value, based on the membership functions of the inputs.

With this basic calculus numerous researchers have shown that imprecision can be handled in a useful and logically consistent manner. In particular, just as ordinary set theory forms the basis for reasoning via logic; fuzzy set theory forms the basis for approximate reasoning, via a fuzzy logic.

The Use of Fuzzy Sets in Managing Expert Information Systems

The main advantages of using fuzzy set theory and fuzzy logic when designing an information system lies in the fact that we can allow fuzzy, or approximate, input from the user. Although it is possible that stochastic, or probability, models might also be appropriate in this situation, using fuzzy logic is faster computationally and thus appears to offer a practicable intelligent system. We have not as yet developed these ideas very far and thus we shall present only an overview of the sorts of ideas that might be worth researching further. It should be noted that other researchers in the field have already begun using these concepts and claim great successes for their work. Such literature as may have been generated, however, does not appear to have been published. The research thus far has been reported only at international congresses, for example the Congress on Cybernetics and Systems in Acapulco, Mexico, December 12 through 17, 1980. The ideas presented in this section, however, do not borrow from those of other researchers, since unfortunately we have not been able to study their work.

An attractive and fairly simple idea would be to extend the use of key words. In order to achieve this, one could envisage associating with each item a value indicating its degree of membership in the fuzzy set of data satisfying a given key word. Such a value could be preprogrammed into the information phase. Then one can envisage using a fuzzy extension of cluster analysis. This mathematical technique has been studied, for example, by Bezdec, and several good results have been reported with it both in terms of intuitiveness of the results and simplicity of the calculation. The different data could be grouped together in fuzzy clusters relating to their zone of applicability and then either the names of the fuzzy clusters could be presented to the data base user, or perhaps the prototype member of the fuzzy clusters could be presented to the data base user. The user could then state which fuzzy clusters he would be interested in viewing. Alternatively, the user could present a prototype to the system which would then present to the user those items of data which were "near" to the prototype, in terms of the fuzzy clustering algorithm. Parsimony of data presentation could be incorporated into the process by allowing only data which had membership greater than a given threshold value in the relevant fuzzy cluster to be presented. Should further data be requested by the user one could then lower the threshold value.

Fuzzy logic is of great potential when contemplating the design of interactive data bases because the approximate concepts in which human beings typically think can be explicitly modeled. For example, the system would understand what was meant by information that was "sort of relevant to Russian submarine maneuvers." With such concepts, a quick and highly interactive search of an unfamiliar data base could be conducted by any user.

It should be realized that the concepts discussed above do not draw on the decision theoretic idea of value of information. Thus

we have done no more than whet the readers appetite for investigating such use of fuzzy sets further.

Fuzzy Decision Analysis and the Value of Coherence

We follow Watson et al. (1979) and Freeling (1980a) and impute from the imprecision in the inputs to a decision model, what the imprecision in the conclusions should be. We assume that the structure of a decision problem is clear-cut, but that the input probabilities and utilities are known only fuzzily (approximately). (We realize that in many decision analyses the appropriate structure is only fuzzily known, but we restrict attention here to problems where this is not the case).

The inputs, then, are fuzzy numbers, as is the output (the expected utility). An example of a fuzzy input and output are input functions shown in Exhibit 1(b) and (c). Note that an ordinary "crisp" probability (1(a)) can have only one value; a fuzzy probability has many possibilities. Such fuzziness may arise because of lack of time in assessments i.e., at a first pass we ask only for approximate values; or because the DM has no more than a vague, or fuzzy, notion in his head of "probability." This motivation is further discussed in Freeling (1980b). We take the view that we may model approximately the results of a more detailed analysis, by assuming that the analysis will increase the DM's coherence by reducing his initial fuzz on the inputs. In particular, we introduce the concept of perfect coherence as the situation when all inputs (and hence outputs) are crisp. In a manner analogous to that used in a normal decision analysis to calculate value of perfect information (VOPI), we may calculate the value of such perfect coherence, before performing the analysis. This, then, may be considered an upper bound on the value of the analysis to the DM--how much he should be willing to pay to achieve increased coherence by decision analysis. One may

MEMBER-
SHIP
FUNCTION

1 +

1

0.5

1

(A)

A CRISP PROBABILITY

MEMBERSHIP
FUNCTION

1 +

0.3 0.5 0.7

1

(B)

A FUZZY PROBABILITY

MEMBERSHIP
FUNCTION

1 +

1

EXPECTED UTILITY

1

(C)

FUZZY EXPECTED UTILITIES


$A = \text{MAX}(A, B)$

EXHIBIT 1



also extend the concept of VOPI to this fuzzy analysis. For the simple decision tree of Exhibit 2, where only the probability p is fuzzy, as shown in (b), the (fuzzy) VOPC and VOPI are shown in (c). It can also be shown that, in a very natural sense, the value of perfect information is always greater than the value of perfect coherence. This satisfying and intuitive result leads us to hope that this approach to the valuation of analysis may prove more fruitful, at least in some situations, than the more traditional ones discussed elsewhere in this report.

We have thus shown that we can, in a logically consistent manner, gain an approximate idea of the value of analysis by asking only for approximate inputs. Such a calculation could, we believe, be fairly simply incorporated into an interactive data base, both to calculate the value of a decision aiding device, and also an approximate value of different types of information.

ACT?



(A)
THE DECISION TREE



(C)
VALUE OF COHERENCE & INFORMATION

DISTRIBUTION LIST

OSD

CDR Paul R. Chatelier
Office of the Deputy Under
Secretary of Defense
OUSDRE (E&LS)
Pentagon, Rm. 3D129
Washington, D.C. 20301

Department of the Navy

Director
Engineering Psychology
Programs
Code 455
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217
(5 copies)

Director
Tactical Development &
Evaluation Support
Code 230
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Naval Analysis Programs
Code 431
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Operations Research Programs
Code 434
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Statistics and
Probability Program
Code 436
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Department of the Navy

Director
Information Systems Program
Code 437
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Special Assistant for Marine
Corps Matters
Code 100M
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Commanding Officer
ONR Eastern/Central Regional
Office
ATTN: Dr. J. Lester
Bldg. 114, Section D
666 Summer Street
Boston, MA 02210

Commanding Officer
ONR Branch Office
ATTN: Dr. C. Davis
536 South Clark Street
Chicago, IL 60605

Commanding Officer
ONR Western Regional Office
ATTN: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Office of Naval Research
Scientific Liaison Group
American Embassy, Room A-407
APO San Francisco, CA 96503

Director
Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D.C. 20375
(6 copies)

Department of the Navy

Dr. Bruce Wald
Communications Sciences Division
Code 7500
Naval Research Laboratory
Washington, D. C. 20350

Dr. Robert G. Smith
Office of the Chief of Naval
Operations, OP987H
Personnel Logistics Plans
Washington, D. C. 20350

Naval Training Equipment Center
ATTN: Technical Library
Orlando, Florida 32813

Human Factors Department
Code N215
Naval Training Equipment Center
Orlando, Florida 32813

Dr. Alfred F. Smode
Training Analysis and
Evaluation Group
Naval Training Equipment Center
Code N-00T
Orlando, Florida 32813

Dr. Arthur Bachrach
Behavioral Sciences Department
Naval Medical Research Institute
Bethesda, Maryland 20014

Dr. George Moeller
Human Factors Engineering Branch
Submarine Medical Research Lab
Naval Submarine Base
Groton, Connecticut 06340

Commanding Officer
Naval Health Research Center
San Diego, California 92152

Dr. James McGrath, Code 311
Navy Personnel Research and
Development Center
San Diego, California 92152

Department of the Navy

Navy Personnel Research and
Development Center
Management Support Department
Code 210
San Diego, California 92152

CDR P. M. Curran
Code 604
Human Factors Engineering Division
Naval Air Development Center
Warmister, Pennsylvania 18974

Mr. Ronald A. Erickson
Human Factors Branch
Code 3194
Naval Weapons Center
China Lake, California 93555

Human Factors Engineering Branch
Code 1226
Pacific Missile Test Center
Point Mugu, California 93042

Dean of the Academic Departments
U.S. Naval Academy
Annapolis, Maryland 21402

Dr. Gary Poock
Operations Research Department
Naval Postgraduate School
Monterey, California 93940

Dean of Research Administration
Naval Postgraduate School
Monterey, California 93940

Mr. Warren Lewis
Human Engineering Branch
Code 8231
Naval Ocean Systems Center
San Diego, California 92152

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
Code RD-1
Washington, D. C. 20380

Department of the Navy

HQS, U.S. Marine Corps
ATTN: CCA40 (MAJOR Pennell)
Washington, D. C. 20380

Mr. Phillip Andrews
Naval Sea Systems Command
NAVSEA 0341
Washington, C. C. 20362

Commander
Naval Electronics Systems Command
Human Factors Engineering Branch
Code 4701
Washington, D. C. 20360

LCDR W. Moroney
Code 55MP
Naval Postgraduate School
Monterey, California 93940

Mr. Merlin Malehorn
Office of the Chief of Naval
Operations (OP 102)
Washington, D. C. 20350

Department of the Army

Mr. J. Barber
HQS, Department of the Army
DAPE-MBR
Washington, D. C. 20310

Dr. Joseph Zeidner
Technical Director
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, Virginia 22333

Director, Organizations and
Systems Research Laboratory
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, Virginia 22333

Technical Director
U.S. Army Human Engineering Labs
Aberdeen Proving Ground, Md. 21005

Department of the Army

U.S. Army Aeromedical Research Lab
ATTN: CPT Gerald P. Krueger
Ft. Rucker, Alabama 36362

ARI Field Unit-USAREUR
ATTN: Library
C/O ODCSPER
HQ USAREUR & 7th Army
APO New York 09403

Department of the Air Force

U.S. Air Force Office of
Scientific Research
Life Sciences Directorate, NL
Bolling Air Force Base
Washington, D. C. 20332

Dr. Donald A. Topmiller
Chief, Systems Engineering Branch
Human Engineering Division
USAF AMRL/HES
Wright-Patterson AFB, Ohio 45433

Air University Library
Maxwell Air Force Base, AL 36112

Dr. Gordon Eckstrand
AFHRL/ASM
Wright-Patterson AFB, Ohio 45433

Foreign Addresses

North East London Polytechnic
The Charles Myers Library
Livingstone Road
Stratford
London E15 2LJ
ENGLAND

Prof. Dr. Carl Graf Hoyos
Institute for Psychology
Technical University
8000 Munich
Arcisstr 21
FEDERAL REPUBLIC OF GERMANY

Other Organizations

Dr. Gershon Weltman
Perceptronica
6271 Variel Avenue
Woodland Hills, California, 91364

Dr. Meredith P. Crawford
American Psychological Association
Office of Educational Affairs
1200 - 17th Street, N.W.
Washington, D. C. 20036

Dr. Ward Edwards
Director, Social Science
Research Institute
University of Southern California
Los Angeles, California 90007

Dr. Charles Gettys
Department of Psychology
University of Oklahoma
455 West Lindsey
Norman, Oklahoma 73069

Dr. Kenneth Hammond
Institute of Behavioral Science
University of Colorado
Room 201
Boulder, Colorado 80309

Dr. Ronald Howard
Department of Engineering-
Economic Systems
Stanford University
Stanford, California 94305

Dr. William Howell
Department of Psychology
Rice University
Houston, Texas 77001

Journal Supplement Abstract
Service
American Psychological Assoc.
1200 - 17th Street, N.W.
Washington, D.C. 20036 (3 copies)

Mr. Richard J. Heuer, Jr.
27585 Via Sereno
Carmel, California 93923

Other Organizations

Mr. Tim Gilbert
The MITRE Corporation
1820 Dolly Madison Blvd.
McLean, Virginia 22102

Dr. John Payne
Duke University
Graduate School of Business
Administration
Durham, North Carolina 27706

Dr. Andrew P. Sage
University of Virginia
School of Engineering and
Applied Science
Charlottesville, Virginia 22901

Foreign Addresses

Dr. Kenneth Gardner
Applied Psychology Unit
Admiralty Marine Technology
Establishment
Teddington, Middlesex TW11 OLN
ENGLAND

Director, Human Factors Wing
Defense and Civil Institute of
Environmental Medicine
P.O. Box 2000
Downsview, Ontario M3M 3B9
CANADA

Dr. A. D. Baddeley
Director, Applied Psychology Unit
Medical Research Council
15 Chaucer Road
Cambridge, CB2 2EF
ENGLAND

Other Governmental Agencies

Defense Documentation Center
Cameron Station, Bldg. 5
Alexandria, Virginia 22314
(12 copies)

Dr. Craig Fields, Director
Cybernetics Technology Office
Defense Advanced Research
Projects Agency
1400 Wilson Blvd.
Arlington, Virginia 22209

Dr. Judith Daly
Cybernetics Technology Office
Defense Advanced Research
Projects Agency
1400 Wilson Blvd.
Arlington, Virginia 22209

Professor Douglas E. Hunter
Defense Intelligence School
Washington, D. C. 20374

Other Organizations

Dr. Robert R. Mackie
Human Factors Research
5775 Dawson Avenue
Goleta, California 93017

Other Organizations

Dr. Gary McClelland
Institute of Behavioral
Sciences
University of Colorado
Boulder, Colorado 80309

Human Resources Research
Office
300 N. Washington Street
Alexandria, Virginia 22314

Dr. Miley Merkhofer
Stanford Research Institute
Decision Analysis Group
Menlo Park, California 94025

Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army-Navy Drive
Arlington, Virginia 22202

Prof. Judea Pearl
Engineering Systems Dept.
University of California
405 Hilgard Ave.
Los Angeles, California 90024

Prof. Howard Raiffa
Graduate School of Business
Administration
Harvard University
Soldiers Field Road
Boston, Massachusetts 02163

Dr. T. B. Sheridan
Dept. of Mechanical Engineer-
ing
MIT
Cambridge, Massachusetts 02139

Dr. Arthur I. Siegel
Applied Psychological Services
404 East Lancaster Street
Wayne, Pennsylvania 19087

Dr. Amos Tversky
Department of Psychology
Stanford University
Stanford, California 94305